LQCD-ext

Computational Resources for Lattice QCD: 2010–2014

Revised, April 12, 2009

Lattice QCD Executive Committee

R. Brower, (Boston U.) N. Christ (Columbia U.), M. Creutz (BNL),
P. Mackenzie (Fermilab), J. Negele (MIT), C. Rebbi (Boston U.),
D. Richards (TJNAF), S. Sharpe (U. Washington), and R. Sugar (UCSB)

1 Introduction

The Department of Energy (DOE) has supported computational infrastructure for the study of lattice quantum chromodynamics (lattice QCD) for the last seven years. It has funded the development of software through two grants from the Scientific Discovery Through Advanced Computing (SciDAC) Program, the acquisition of clusters optimized for the study of lattice QCD through the first of the SciDAC grants (SciDAC I) and the current Lattice QCD Computing Project (LQCD), and the construction of the specially designed QCDOC computer through a stand alone grant. The operation of the clusters and the QCDOC is supported through LQCD. The software produced under the SciDAC grants is publicly available, and the hardware is open, on a peer reviewed basis, to all members of the USQCD Collaboration, which consists of nearly all of the high energy and nuclear physicists in the United States involved in the numerical study of lattice QCD. The DOE funded infrastructure has enabled major progress in the study of lattice QCD, and helped to bring the field to a point where it is now providing accurate determinations of a wide range of quantities of importance to experimental programs in high energy and nuclear physics.

The current SciDAC grant (SciDAC II), which funds software development, runs through March 14, 2011, whereas LQCD, which funds the acquisition and operation of hardware, ends on September 30, 2009. The committee of scientists that reviewed LQCD in the spring of 2007 stated in its report that "*The resources provided through the LQCD project are crucial for the US lattice QCD community to stay internationally competitive. This will remain true beyond the final year of the LQCD project, 2009, and the committee believes that an increase in computational resources beyond 2009 should be strongly encouraged, building on the success of the 2006–2009 LQCD project.*" We agree, and in this proposal we set out a plan for the acquisition and operation of dedicated hardware for the study of lattice QCD for the fiscal years 2010–2014.

In Section 2 of this document we set out our scientific objectives for the period 2010–2014, and indicate the potential impact of achieving them on DOE's experimental programs in high energy and nuclear physics. Work will focus on precision tests of the Standard Model and determination of its fundamental parameters; the study of strongly interacting matter under extreme conditions of temperature and density; the calculation of the masses, internal structure and interactions of strongly interacting particles; and the exploration of strongly coupled field theories that go beyond the Standard Model. In Section 3 we describe the computational resources needed from the DOE to achieve our scientific goals. They consist of a roughly equal mixture of cycles on the DOE's leadership class computers and on less capable, but still powerful dedicated hardware. The purpose of this proposal is to obtain the funds to acquire and operate the dedicated hardware, which we plan to locate at the three laboratories that have housed the LQCD hardware: Brookhaven National Laboratory (BNL), Fermi National Accelerator Laboratory (FNAL), and Thomas Jefferson National Accelerator Facility (TJNAF). We request \$2.01M (\$2,010,000) per year for hardware acquisition, and an operations budget that starts at \$1.13M in 2010 and rises to \$1.68M in 2014. In Section 4 we discuss the role of the USQCD Collaboration, and explain the processes by which it sets scientific priorities and allocates computational resources. We also describe the international collaborations in which members of USQCD are engaged, the sharing of large data sets through the International Lattice Data Grid (ILDG), and the software created under our SciDAC grants. In Section 5 we discuss the role of the participating laboratories, and in Section 6 we set out the proposed management structure for LQCD-ext

2 Scientific Objectives

The objective of LQCD-ext is to achieve the DOE Office of Science's 2009 strategic milestone to "*Use computer simulations to calculate strong interactions between particles so precisely that theoretical uncertainties no longer limit our understanding of these interactions* ..." The level of accuracy needed to achieve this milestone has been achieved for a limited number of quantities, and tools have been developed to do so for a wide variety of others provided the computational resources requested in this proposal are available. In this Section we describe the four major areas on which work under LQCD-ext will focus, and discuss the relationship of this work to experimental programs in high energy and nuclear physics [1]. As discussed in Section 4.1, a process is in place for ordering scientific priorities and allocating USQCD resources on a yearly basis. Detailed priorities and computational approaches may evolve over time, as new experimental results and calculation methods appear; however, we fully expect that the broad goals and the estimates of the resources needed to achieve them will not change.

2.1 Fundamental Parameters from Future Lattice Calculations

One of the central aims of calculations using lattice QCD is to determine the underlying parameters of the Standard Model (SM) by stripping away the effects of the strong interactions. Lattice calculations aim to provide accurate determinations of the quark masses, the strong coupling constant α_S, and the values of the weak transition couplings between quarks—i.e. the elements of the Cabibbo-Kobayashi-Maskawa (CKM) matrix. Particularly exciting is the possibility of determining different, inconsistent values of the CKM matrix elements from different decay processes. This would indicate a breakdown in the Standard Model and thus the need for new physics. This approach is complementary to the direct discovery searches to be undertaken at the Large Hadron Collider at CERN (LHC). Furthermore, if new physics is found, then the constraints from rare SM processes, many of which require lattice QCD calculations, will be needed, together with LHC results, to work out the details.

The last five years has seen very significant progress in the lattice QCD calculations needed to determine the SM parameters. Accurate calculations of many quantities are now possible, with all errors controlled. This has allowed extensive validation of the method by comparing lattice and experimental results [2, 3, 4], including several successful predictions [5]. Crucial to this success has been the creation, by the USQCD collaboration, of an ensemble of gauge configurations generated including the full quantum measure. These "unquenched" configurations include the effects of light quark loops, with a series of values for the lattice spacing (a) and the light quark masses.[1] This ensemble allows one to do controlled chiral and continuum extrapolations. The importance of using unquenched configurations, and the present level of validation, is shown in Fig. 1. The agreement of all unquenched lattice results having percent-level accuracy with experiment gives us confidence in both heavy and light quark methodologies.

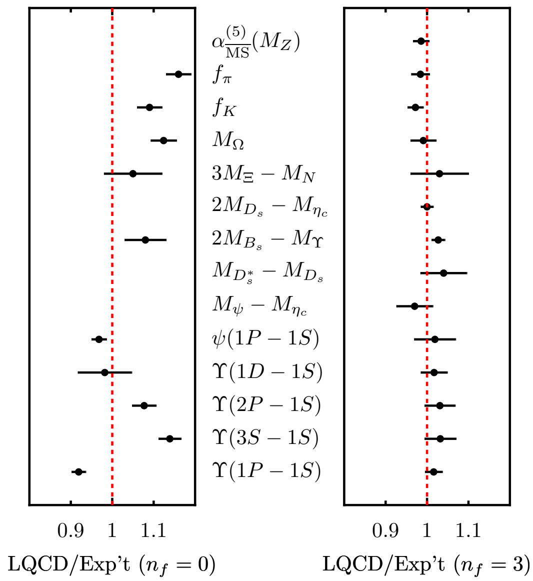

Figure 1: Ratio of accurately determined lattice results (for so-called "gold plated" quantities) to those from experiment. Left panel uses quenched and right panel unquenched ($2+1$ flavor) lattice calculations. The overall scale is determined in both cases from the $\Upsilon(1P) - \Upsilon(1S)$ splitting, while quark masses are fixed from m_π, m_K, $m_{J/\psi}$ and the Υ spectrum. From Ref. [6].

These calculations also give controlled determinations of the quark masses and α_S (the latter shown in Fig. 1, compared to results from matching high-energy experiments to perturbative QCD). Although the methodology is different for light and heavy quarks, it turns out that present errors in both are comparable, and of size 7-10% [7, 8, 9].[2] Quark masses are not only of interest as fundamental parameters of the SM, but also are required as inputs into predictions of rare decays. For example, m_c is needed to predict the SM contribution to $K \to \pi\nu\bar{\nu}$ (and thus constrain the contributions of new physics) while m_π^2/m_ℓ is needed as input into predictions from the Soft-Colinear effective theory for non-leptonic B decays. Thus we intend to systematically reduce the errors in these masses, and expect percent-level control by the end of the LQCD-ext.

[1] Calculations to date are in the isospin symmetric limit $m_u = m_d = m_\ell$, but have the strange quark mass, m_s, at (or close to) its physical value. The simulated values of m_ℓ are larger than the physical average light quark mass, but range down to $m_\ell/m_s \approx 1/10$, which approaches the physical ratio of $\approx 1/27$.

[2] In detail: $m_u = 2.0(2)\,\text{MeV}$, $m_d = 4.6(3)\,\text{MeV}$, $m_s = 90(7)\,\text{MeV}$, $m_c = 1.22(9)\,\text{GeV}$ and $m_b = 4.4(3)\,\text{GeV}$.

Quantity	CKM element	present expt. error	present lattice error	2009 lattice error	2014 lattice error	error from non-lattice method
f_K/f_π	V_{us}	0.3%	0.9%	0.5 %	0.3%	–
$f_{K\pi}(0)$	V_{us}	0.4%	0.5%	0.3%	0.2%	1% (ChPT)
$D \to \pi\ell\nu$	V_{cd}	3%	11%	6%	4%	–
$D \to K\ell\nu$	V_{cs}	1%	11%	5%	2%	5% (ν scatt.)
$B \to D^*\ell\nu$	V_{cb}	1.8%	2.4%	1.6%	0.8%	$< 2\%$ (Incl. $b \to c$)
$B \to \pi\ell\nu$	V_{ub}	3.2%	14%	10%	4%	10% (Incl. $b \to u$)

Table 1: Present status and future prospects for lattice calculations which directly determine elements of the CKM matrix. All errors are quoted for the CKM elements themselves. Estimates are from the contributions of Juettner, Laiho, Lüth, Shipsey, and van de Water to Ref. [10], and from Ref. [11]. The last column, if present, shows the present error attainable on the CKM element using competing, non-lattice approaches.

The work of the last five years has set the stage for fully controlled calculations of the more complicated matrix elements needed to determine or constrain CKM elements. Indeed, for the simplest such quantities—single-particle matrix elements involving vector and axial currents—first unquenched results are already available. By calculating semi-leptonic form factors (and also f_K/f_π), one can determine all CKM elements which do not involve the top quark, as shown in Table 1.[3] These results represent a major milestone for the field. Nevertheless, comparing the third and fourth columns one sees that errors from lattice simulations exceed those from experiment for all quantities. Lattice errors are also not yet competitive with those from other methods for V_{cs} and V_{ub}.

An important goal of our future program is to reduce these lattice errors so that, ultimately, they are smaller than the experimental errors.[4] The table gives estimates of what should be possible in LQCD (2009 error) and a fully funded LQCD-ext (2014 error). We see that LQCD will not attain the desired goal, whereas LQCD-ext will approach it—giving lattice errors comparable to, or smaller than, experimental errors for most of the quantities. The precision attainable with LQCD-ext is important for several reasons. First, it will allow much more stringent tests of the unitarity of the CKM matrix. For example, using the present lattice V_{us} from f_K/f_π, one has

$$|V_{ud}|^2 + |V_{us}|^2 + |V_{ub}|^2 = 0.9977(5)(12)(0)\,, \tag{1}$$

where the errors are respectively from those in V_{ud}, V_{us} and V_{ub}. While this is consistent with unitarity, reducing the lattice error by a factor of 3, which we estimate to be possible in LQCD-ext, would provide a significantly more stringent test. Accurate CKM elements are also crucial for using several rare SM decays to constrain new physics. For example, the SM rate for $K \to \pi\nu\overline{\nu}$ decays is proportional to V_{cb}^4. Finally, precision lattice results will allow one to determine the significance of the present marginal disagreement between V_{ub} determined from exclusive $B \to \pi$ decays (using lattice input) and from inclusive $b \to u$ decays (using HQET).

The remaining elements of the CKM matrix involve the top quark, and are thus accessible only through second-order weak processes involving top quark loops. These processes are particularly

[3] V_{ud} is not included, since its determination using the lattice f_π and the $\pi \to \mu\nu$ decay rate is not competitive with that from nuclear $\beta-$decays. The lattice result for f_π is used instead as a validation, as shown in Fig. 1.

[4] Experimental errors will themselves be reduced for most of these quantities, providing even stronger motivation for improved lattice calculations.

interesting since the contributions of new physics can more easily appear above the "background" of the small SM rate, assuming that the latter can be calculated accurately. Furthermore, in three examples—$B_d - \overline{B}_d$ and $B_s - \overline{B}_s$ mixing, and CP violation in $K - \overline{K}$ mixing—accurate experimental results are available and precision lattice results are possible. The lattice must provide matrix elements of four-fermion operators, which are more challenging than those of currents just discussed, so that results lag somewhat behind. Unquenched results at only a single lattice spacing are available to date, but results with all errors controlled should be announced in 2008. The present status and future prospects are collected in Table 2.

Hadronic Matrix Element	UT result current	Lattice errors current	Lattice errors 2009	Lattice errors 2014
$\widehat{B}_K$	0.78 ± 0.09	0.77 ± 0.05	± 0.025	± 0.01
$f_{B_s}\sqrt{\widehat{B}_{B_s}}$	261 ± 6 MeV	282 ± 21 MeV	± 10 MeV	± 5 MeV
ξ	1.25 ± 0.06	1.23 ± 0.06	± 0.02	± 0.01

Table 2: Status of and prospects for lattice calculations of matrix elements which play a key role in the determination of CKM matrix elements. UT results from Ref. [12] are explained in the text. Present estimates are from Refs. [13] (B_K), [14] ($f_{B_s}\sqrt{B_{B_s}}$), and [15] ($\xi = f_{B_s}\sqrt{B_{B_s}}/(f_B\sqrt{B_B})$). Future error estimates are taken from the estimates of Juettner and Gamiz to Ref. [10] and from Ref. [11].

Reducing errors to the expected 2-4% level by the end of LQCD will be a major milestone. To test reliability we will obtain results from at least two different methods in each case, e.g. with domain wall fermions (DWF) and mixed DWF/staggered fermions for B_K, and with bottom quarks discretized using both the Fermilab formalism and NRQCD. More extensive ensembles, with smaller a and m_ℓ and larger volumes, will also be needed. Our plans for these ensembles are described below.

The importance of obtaining precision lattice results has been increased by the enormous progress in measuring B-meson properties in the last five years. Accurate results for CP-violation in several decays has allowed a determination of CKM elements, and in particular the "unitarity triangle" (UT), with little input from hadronic physics (and with no lattice input). The resulting constraints on the Wolfenstein parameters $\bar{\rho}$ and $\bar{\eta}$ are shown in the left panel of Fig. 2. Lattice results for the matrix elements in Table 2 allow an independent determination of $\bar{\rho}$ and $\bar{\eta}$, based on different experiments, with the present status shown in the right panel. The SM requires that these two constraints must be consistent—as they are at present. Assuming this consistency, i.e. that the SM is correct, one can turn the results in the left panel into "predictions" of the hadronic matrix elements. These are given in the "UT results" column of Table 2.

In the next 5-7 years, the errors in these UT results are likely to decrease substantially, mainly due to improved determinations of the CKM elements of Table 1 (based on both experimental and lattice improvements). The net result is that errors in the matrix elements of Table 2 at the 1-2% level will be needed to keep pace. We estimate that such errors will be attainable by the end of LQCD-ext. This will require the use of new methods—non-perturbative matching for all quantities, and possibly discretized HQET. Our goal is to constrain the SM at the 1% level in a sector where new physics is most likely to appear.

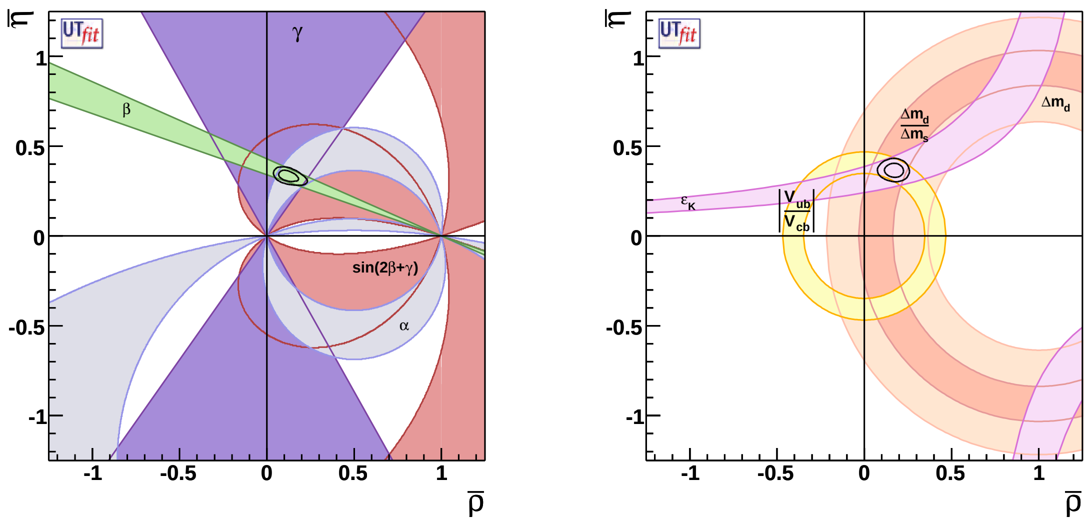

Figure 2: Present constraints on $\bar\rho$ and $\bar\eta$ from the UT analysis (left panel) and from matrix elements involving lattice QCD input (right panel). Contours of 68% and 95% probability are shown, together with the 95% probability regions from individual constraints. From Ref. [12].

In addition to the matrix elements of Tables 1 and 2, there are many others that the lattice can calculate. For the sake of brevity we simply list some important examples—for more details see Ref. [11]. Several quantities will provide further validation: f_D, f_{D_s}, charmonium spectroscopy, spectroscopy of excited D and B mesons, c and b baryon spectroscopy, hyperon semi-leptonic decays. Other quantities provide further tests of the SM: $(\Delta\Gamma/\Gamma)_{B_s}$, $B \to K\ell^+\ell^-$ form factors. Finally, there are matrix elements of operators which appear in beyond the SM (BSM) theories but which are absent in the SM. These are thus needed to constrain the parameters of particular BSM theories, so as to ensure that they are consistent with rare SM processes such as meson-antimeson mixing. Examples include $B-\overline{B}$, $D-\overline{D}$ and $K-\overline{K}$ mixing amplitudes from non-SM four-fermion operators, proton decay matrix elements and hadronic contributions to $g-2$. For all these quantities the calculations are of a similar level of difficulty to those described above, and thus precision results should be attainable. The desired precision depends on the accuracy of the corresponding experimental results. We expect to continuously add new quantities to our repertoire so that those listed above, and others, will be available during LQCD-ext.

Thus far we have described the impact that LQCD-ext can make largely by extending existing methods. We now turn to calculations that LQCD-ext will allow that are not possible otherwise. These involve either two particle states (or resonances) and/or quark disconnected diagrams. The former require larger volumes than studied to date ($L = 5-6$ fm as opposed to $2.5-3$ fm) as well as using finite volume corrections in an essential way. The latter require new techniques for noise reduction. There has been much effort devoted in the lattice community to developing the necessary methods, and we are confident that LQCD-ext will allow substantial progress. We are unable, however, to provide the precise error estimates given above as there are no available calculations from which to extrapolate.

One of the most important challenges faced by lattice QCD, and a major focus of this proposal, is the calculation from first principles of direct CP violation in two-pion decay of the neutral kaon. This is described by the ratio $\varepsilon'/\varepsilon = (16.5 \pm 2.6) \times 10^{-4}$, which was measured at both CERN [16] and Fermilab [17]. Because the kaon mass lies below the four-pion threshold and the strong interactions conserve G parity (preventing mixing between two and three pion states), finite volume methods [18] in principle permit the direct calculation of the relevant $K \to \pi\pi$ decays. However, such calculations have not been possible to date given the large volumes which are required, $L \approx 6\,\mathrm{fm}$.

In addition, much can be learned about these processes using chiral perturbation theory. At leading order, one need only calculate the more tractable $K \to \pi$ and $K \to 0$ amplitudes, although these involve quark disconnected loops. Following the proof-of-principle quenched studies using DWF, an unquenched calculation using this approach is underway [19]. It is expected that these calculations, to be completed in 2008, will be limited by the large violation of chiral symmetry that occurs for masses as large as m_K and will yield errors on the order of 30%. Over the next couple of years, this work will be extended to the full $K\pi\pi$ vertex to provide additional constraints on the chiral perturbation theory description [20, 21, 22].

The larger volumes and smaller masses targeted by this proposal will permit further refinement of the study of $K \to \pi\pi$ amplitudes using chiral perturbation theory. What is much more important, however, is that these larger volumes and smaller quark masses will permit direct calculation of $K \to \pi\pi$ decay amplitudes as matrix elements between the kaon state and proper, energy conserving two-pion final states. Earlier calculations [23] demonstrate that with these volumes and quark masses calculations of the $\Delta I = 3/2$ amplitudes should be possible with errors well below 10%. The more important $\Delta I = 1/2$ are much more difficult due to operator mixing and quark-disconnected contributions. However, errors on the order of 20% may be possible using either twisted boundary conditions [24] or non-zero center of mass momentum [25]. Considerable effort will be invested in advance of 2010 to prepare for this opportunity.

Other quantities that can be calculated on large volumes are amplitudes involving resonances. Of particular interest for constraining the SM are $B \to K^*\gamma$, $B \to \rho\ell\nu$ and $B \to \rho\ell\ell$ form factors.

Once calculations involving disconnected quarks are possible, there are many additional quantities of interest. A partial list is the nucleon matrix elements of $\bar{q}q$, which control the sensitivity of some dark-matter detectors, $B \to \eta\ell\nu$ form factors, the neutron electric dipole moment, glueball-$\bar{q}q$ mixing, and lifetime ratios for bottom hadrons.

We close this section with a description of possible ensembles of configurations. For the quantities for which we are aiming for percent-level or smaller precision, our main challenge is to reduce the systematic errors from the continuum, chiral and infinite volume extrapolations, as well as from matching of continuum and lattice operators. Thus, on the numerical side, we need to reduce the lattice spacing and the minimal light-quark mass, while holding the volume large enough that finite volume errors are smaller than our desired precision. In addition, we need large enough ensembles that statistical errors are smaller than those from systematics. Finally, we need to improve the matching calculations— from one-loop to either two-loop or non-perturbative matching.

Our present work concerning fundamental parameters uses lattices generated with both improved staggered fermions and DWF, and we assume here that we will continue with this combination. We stress, however, that the choice of fermion action will be decided by the collaboration through the mechanism described in Section 4.1, based on its estimate of the optimal approach to obtaining timely physics results. Our simulations with staggered fermions use a rooted determinant to deal with the effects of the additional fermion varieties (”tastes”) that are intrinsic to this formulation. Previous studies have shown that many quantities of interest can be calculated to high accuracy with this approach [2, 3, 4, 7].

We now give examples of what will be possible given the computational power we would obtain with the proposed LQCD-ext resources together with estimated time on leadership class machines. We expect that the generation of configurations and some general-purpose propagators will be done on the leadership-class machines, while the calculation of most of the propagators together with the analysis will be done on USQCD dedicated hardware. We expect a ratio of about 1:1 for these two components. Thus it is not unreasonable that 200 Tflop-Years [26] will be available for configuration generation for flavor physics in LQCD-ext. Given our expectation that we can roughly double our resources by sharing configurations with researchers in other countries, we give an example of lattices that will require about 400 Tflop-Years in total. We note that all estimates are for sustained performance [27].

a (fm)	m_ℓ/m_s	Size	m_π (MeV)	L (fm)	Lm_π	MC traj.	TF-Yrs
0.09	0.075	$48^3\times 96$	200	4.3	4.4	5000	0.5
0.06	0.10–0.4	$72^3\times 144$	230	4.3	5.0	5000	7.1
0.06	0.075	$72^3\times 144$	200	4.3	4.4	6000	4.3
0.06	0.050	$72^3\times 144$	163	4.3	3.6	6000	7.0
0.06	TOTAL						18.4
0.045	0.15–0.4	$96^3\times 192$	282	4.3	6.2	5000	18.9
0.045	0.1	$96^3\times 192$	230	4.3	5.0	6000	13.5
0.045	0.075	$96^3\times 192$	200	4.3	4.4	6000	19.1
0.045	0.05	$96^3\times 192$	163	4.3	3.6	6000	32.1
0.045	TOTAL						83.6

Table 3: CPU requirements in Tflop-Years (TF-Yrs) for possible future unquenched configurations with HISQ fermions. “MC traj.” gives the lengths of the runs in molecular dynamics trajectories. These lengths are chosen so that statistical errors should be sub-dominant for quantities of interest. The first $a = 0.06$ fm row gives the total Tflop-Years for ensembles with this lattice spacing and $m_l/m_s = 0.1$, 0.15, 0.2 and 0.4. The values of m_π and Lm_pi is this row are for $m_l/m_s = 0.1$. Similarly, the first $a = 0.045$ fm row gives the total TF-Yrs for ensembles with $m_l/m_s = 0.15$, 0.2 and 0.4. The values of m_π and Lm_pi is this row are for $m_l/m_s = 0.15$. All estimates are for two degenerate light quarks of mass m_l and a strange quark at its physical mass.

We expect that ensembles generated with staggered fermions will use the “highly improved staggered quark” or HISQ action [28]. This reduces discretization errors by a factor of 2–3 compared to the second-generation “Asqtad” action which we currently use. By the end of LQCD we will have completed a substantial ensemble with Asqtad sea-quarks: $a = 0.15, 0.12, 0.09, 0.06$ fm with m_ℓ/m_s down to 0.1 and $Lm_\pi \geq 3.5$. This ensemble is the basis for many of the 2009 estimates given above. We also expect within LQCD to produce a HISQ ensemble of similar size to the present MILC ensemble, with $a = 0.15, 0.12, 0.09$ fm, $m_\ell/m_s \geq 0.1$ and $Lm_\pi \geq 3.6$. This will take about 1.3 Tflop-Years in total, and will allow us to validate this action and optimize our generation codes. We note that running at $a = 0.09$ fm with the HISQ action leads to similar discretization effects to running at $a = 0.06$ fm with the asqtad action. A possible plan for HISQ running in LQCD-ext is presented in Table 3. We aim to push two further steps downwards in lattice spacing

a (fm)	m_ℓ/m_s	Size	m_π(GeV)	L (fm)	Lm_π	MC traj.	TF-Yrs
0.086	0.150	$48^3\times64\times16$	276	4.1	5.8	6000	7
0.086	0.093	$48^3\times64\times16$	217	4.1	4.5	2400	3
0.086	TOTAL						10
0.125	0.134	$32^3\times64\times24$	250	4.0	5.1	4000	2
0.125	0.102	$32^3\times64\times24$	218	4.0	4.4	4000	2
0.125	0.102	$48^3\times64\times24$	218	6.0	6.6	4500	9
0.125	0.071	$48^3\times96\times24$	181	6.0	5.5	4500	18
0.125	0.039	$48^3\times96\times24$	135	6.0	4.1	6000	34
0.125	TOTAL						65
0.094	0.102	$64^3\times96\times24$	218	6.0	6.6	6700	65
0.094	0.071	$64^3\times96\times24$	181	6.0	5.5	7000	78
0.094	0.039	$64^3\times96\times24$	135	6.0	4.1	7600	115
0.094	TOTAL						258

Table 4: CPU requirements for possible future unquenched ensembles using DWF. Lattices are five dimensional, with the last entry in the "size" being the length of the fifth dimension. An improved action will be used for the simulations at $a = 0.125$ and 0.094 fm, as discussed in the text. We estimate the (at present unknown) extra computational cost of simulating with the improved action by increasing the fifth dimension by 50%.

(each roughly halving a^2), and to halve m_ℓ. We choose the same volume for all lattices because we expect to perform partially quenched calculating using the same valence quark masses on each of them (running down to $m_\ell/m_s \approx 0.05$), and thus wish to maintain small finite volume errors from pions composed of valence quarks.

Our tentative proposal for DWF, shown in Table 4, addresses two goals. The first is to complete the set of lattice ensembles begun in 2005 with $a = 0.11$ and 0.086 fm and the Iwasaki gauge action. The plan is to extend the collection of ensembles being created under LQCD by increasing the statistics for the smallest light quark mass ensemble and by adding a larger volume ensemble at the next-to-lightest quark mass. We anticipate completing this goal in 2010.

Our second goal is to move rapidly toward physical light quark masses and correspondingly larger physical volumes. This allows us to use the advantage of DWF—good chiral symmetry—to its full extent. In order to reach substantially lighter quark masses we must decrease the residual quark mass resulting from the small chiral symmetry breaking that is present for DWF when the fifth dimension (L_s) is finite. This will be achieved using an improved action (likely a combination of using a twisted-mass Wilson determinant ratio [29, 30] [with non-zero imaginary mass in the numerator so that topology is sampled ergodically], or Moebius fermions [31], or both). We aim to reduce the residual mass to 0.75 MeV (compared to 5 and 1 MeV on present coarse and fine lattices, respectively), so that physical quark masses can be reached with the contribution from residual chiral symmetry breaking being 30% or less.

Once we change to an improved action the discretization errors are altered, and there is no advantage to continuing with the present lattice spacings (0.11 and 0.086 fm). Thus we propose to use two somewhat coarser ones, 0.125 and 0.094 fm, as shown in the table. This allows us to reach the large volumes and small quark masses required for the study of $K \to \pi\pi$ decays within LQCD-ext.

Our tentative plan is to first undertake the simpler coarse lattice calculations for a sequence of light quark masses running down to the physical value. This would eliminate the chiral extrapolation error and represent a watershed for lattice QCD calculations. The volumes are chosen to keep $L m_\pi \geq 4$. We would then repeat the calculations at the finer lattice spacing, maintaining the same physical volume and the quark masses, so as to determine the size of lattice spacing artifacts.

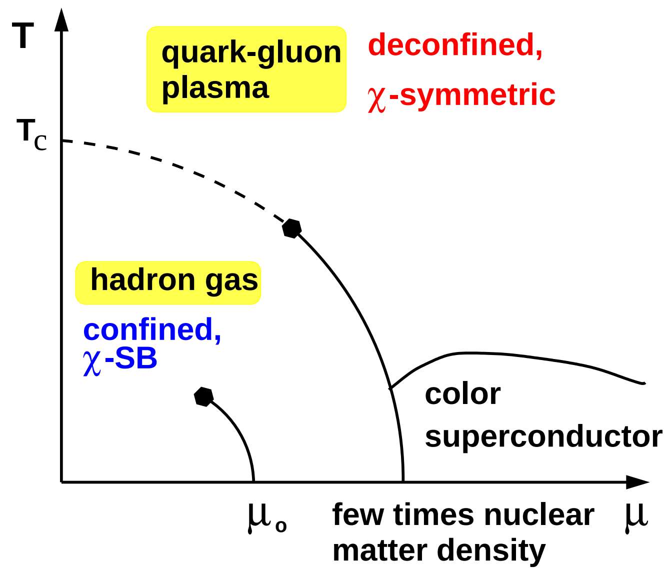

Figure 3: Phase diagram of strongly interacting matter.

2.2 QCD at Nonzero Temperature and Density

Under extreme conditions of high temperature or high baryon number density, strongly interacting matter is expected to have a rich phase structure as indicated in Fig. 3. Characterizing and quantifying the drastic changes in the interaction among elementary particles that go along with such phase changes requires a vigorous experimental program accompanied by large scale numerical calculation.

Over the next decade high performance computing resources will reach the petaflops scale. Coupled with current and planned experiments at the Relativistic Heavy Ion Collider (RHIC), the European heavy ion facility (FAIR), and the Large Hadron Collider (LHC), these computational resources will offer significant opportunities for the advancement of our understanding of the properties of strongly interacting matter at high temperatures and densities.

Lattice quantum chromodynamics is our only source of *ab initio* information about the properties of strongly interacting matter at or near thermal equilibrium and at or near zero baryon number density. Phenomenological models, e.g. hydrodynamics, extend beyond equilibrium and provide the bridge between lattice QCD and the experimental regime. Thus lattice QCD simulations enlarge our understanding of the properties of matter at high temperatures and constrain and validate the phenomenological models.

In the following sections we identify four key computational projects requiring approximately 100 Tflop-Years [26] each that promise significant quantitative and qualitative gains in our knowledge of (1) the equation of state at zero and nonzero density, (2) plasma structure and transport properties, and (3) the phase diagram of QCD at zero and nonzero density.

We estimate the computational cost of each project in the subsections below. The estimates are summarized in Table 5. A brief explanation of some of the simulation parameters is in order. The simulation temperature is determined from a, the lattice spacing, and N_τ, the extent of the lattice in the Euclidean time direction, according to $T = 1/(N_\tau a)$. Thus at any given temperature, the approach to the continuum requires a larger N_τ.

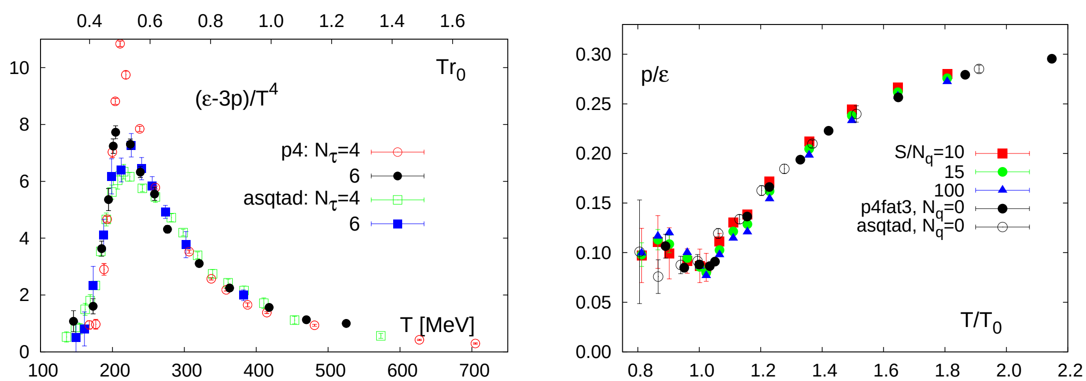

Figure 4: Results for the EoS based on recent calculations with improved staggered quarks. On the left we show the difference of energy density and three times the pressure [32, 33], which is sometimes called the interaction measure. It summarizes our current knowledge of the EoS at low and high temperature at vanishing chemical potential. On the right we compare the ratio of pressure (p) to energy density (ε) at zero baryon number and at nonzero baryon number along curves of fixed entropy (S) per quark (or baryon) ($N_B = N_q/3$) [34]. Here the temperature is given in units of the crossover temperature T_0.

2.2.1 Quark-gluon plasma equation of state

In the hydrodynamic modeling of the expansion of dense matter created in a heavy ion collision, an accurate determination of the equation of state and its associated thermal quantities, namely, energy and entropy density, pressure, and the velocity of sound, are of central importance.

Lattice methods for determining the equation of state (EoS) are well developed, but numerically intensive. Figure 4 illustrates our present knowledge of the continuum EoS. It comes with statistical errors of order 15% and probably comparable systematic errors. Statistical errors can be decreased with longer runs. The principle sources of systematic error are lattice artifacts that decrease with decreasing lattice spacing.

The most successful formulation of lattice quarks for QCD thermodynamics, staggered quarks, has extra quark species that give rise to “taste” multiplets of mesons, including pions. Although we can compensate through standard methods for the additional degrees of freedom, at the coarse lattice spacing of current simulations a pion gas is not accurately simulated. When the lightest pion has the correct experimental mass, the other pion species in the taste multiplet are heavier.

The possible consequences for the equation of state at temperatures slightly below the crossover are not well understood.

The promising, but considerably more computationally expensive DWF formulation of lattice quarks, has the correct number of mesons from the start, but has a related lattice artifact, the "residual quark mass", that gives the pion an unwanted additional effective mass. This artifact can also be systematically and indefinitely reduced, but at further computational expense.

There are two ways to reduce lattice artifacts: reduce the lattice spacing a and improve the action. No matter the quark formulation, the computational cost of an equation of state study grows roughly as a^{-11} with quark masses fixed in physical units. Thus the greatest advances over the past decade have come from improvements in the action that reduce lattice artifacts. Obviously, this is the approach we embrace. In particular, the lattice community has developed increasingly improved formulations of staggered quarks that reduce dramatically the splitting of the taste multiplets. The most successful to date is the highly improved staggered quark (HISQ) action [35]. Compared with the extensively studied "asqtad" action, the HISQ action reduces taste splitting in the pion multiplet by over a factor of two at an additional computational cost of only 30 - 40%. We are currently evaluating the suitability of this action for use in thermodynamic studies.

We plan a calculation of the equation of state with a lattice extent of $N_\tau = 12$ for temperatures up to two times the transition temperature. The lattice spacing in the transition region is approximately 0.09 fm. Lattice artifacts would be reduced by more than half from present simulations, allowing substantial reductions in the error in the extrapolation to zero lattice spacing. The bulk thermodynamic quantities of the equation of state are renormalized by subtracting their values at zero temperature. Thus the determination of the equation of state requires a series of simulations at different temperatures (lattice size $48^3 \times 12$) together with a series of matched zero-temperature simulations (lattice size 48^4) at the same lattice parameters. We estimate the computational cost of the project to be 110 Tflop-Years.

2.2.2 Domain Wall Fermion cross check

The domain wall formulation does not have the spectral lattice artifacts of staggered quarks. However, it is far more computationally expensive. The quality of its chiral behavior is measured by the "residual quark mass" parameter. At $N_\tau = 10$ in the transition region, we know from other tests that this quantity is small enough to assure good control of the residual mass. Our goal is to check the determination of the crossover temperature T_c and the strength of the peak in the chiral susceptibility. We intend to perform a calculation at four temperatures in this critical region on a $48^3 \times 10 \times L_s$ lattice with the fifth-dimension, $L_s = 96$. A calculation at twice the physical light quark mass would be sufficient for a comparison with staggered quark calculations and will allow us to judge systematic effects that may arise in thermodynamic studies through violations of chiral symmetry. We estimate the computational cost of the project to be 80 Tflop-Years.

2.2.3 Equation of state at nonzero density

Heavy ion collisions occur in a baryon-rich environment, whereas lattice simulations are naturally suited for zero baryon density, *i.e.* zero baryon chemical potential. For technical reasons direct simulation at nonzero density and appropriately large lattice volume is extremely difficult. To reach a small, nonzero baryon number density, one constructs the Taylor series expansion in the chemical potential [36]. The coefficients of the series are evaluated in a standard simulation at zero chemical potential. This method is effective for the relatively low baryon number densities of heavy ion collisions. As more terms in the Taylor are calculated, it becomes possible to push to higher chemical potential. The state of the art of such calculations is given in Fig. 4 (right).

Upcoming low energy runs at RHIC will achieve higher baryon density. Present determinations lose statistical significance beyond sixth order in the Taylor series. Knowing higher terms allows us to extend predictions to higher baryon density. We propose a calculation of the equation of state at nonzero density by means of a determination of Taylor series coefficients up to eighth order.

The calculation reuses the lattices generated in the equation of state study. It involves a large set of inversions of the fermion matrix with different random starting vectors. The number of random vectors needed to reach comparable statistical error in different orders N of the Taylor expansion grows exponentially, *i.e.* roughly as 4^N. Moreover, the computational effort per set of random vectors increases approximately as 1.5^N. The overall computational effort thus rises like 6^N. Extrapolating in this way from present calculations, we estimate the cost of this project to be 80 Tflop-Years.

2.2.4 Plasma Structure and Transport Properties

Deconfinement implies the dissolution of hadrons into their constituents. Thus, one would expect that an experimental signal for deconfinement in heavy ion collisions is the disappearance of the charmonium and bottomonium peaks in dilepton production. Lattice simulations and the analysis of experimental measurements suggest, however, that hadronic matter is strongly interacting at temperatures well above T_c. Consequently quarkonium production is suppressed to a degree that depends on temperature. Except for some exploratory studies with dynamical staggered quarks [37] quarkonium spectroscopy at high temperature so far has been done in quenched lattice calculations [38, 39]. As the calculation of heavy quark correlation functions is computationally not very demanding, their analysis in QCD with light quarks will naturally be a part of the studies of transport properties in the light quark sector which we describe in the next paragraph, at almost no extra cost. However, it still is important to complete the current studies in the quenched approximation through investigations of the momentum dependence of quarkonium suppression, which requires large spatial lattice size in order to reach non-zero momenta that are smaller than the temperature. We estimate the cost of this project to be 20 Tflop-Years.

A major puzzle from RHIC experiments is the large degree of collectivity, *i.e.* large elliptic flow. If this happens because the system is thermalized to a good approximation, the degree of quarkonium suppression should provide an estimate of the temperature. So far, relatively little is known about quarkonium properties at nonzero temperature; therefore, lattice information is crucial. If the system is locally thermalized, it should have a very low shear viscosity, *i.e* very small mean free path to produce the observed flow. On the other hand, the bulk viscosity is expected to rise dramatically in the vicinity of the QCD transition. This will be even more relevant in the vicinity of a second order transition point, as it might exist at non-zero density. The bulk viscosity is expected to diverge at a second order transition point.

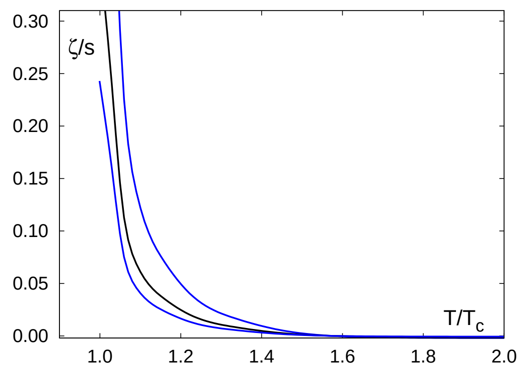

Figure 5: The ratio of bulk viscosity and entropy density determined from QCD sum rules using lattice results on the QCD equation of state as input [40].

While the analysis of heavy quark properties requires the calculation of hadron correlation functions, one needs to analyze correlation functions of the energy momentum tensor to extract viscosities [40, 41]. We show in Fig. 5 a recent estimate of the bulk viscosity based on lattice studies of the trace anomaly (energy density minus three times the pressure) in QCD on lattices with temporal extent $N_\tau = 6$ and phenomenological input from QCD sum rules [40]. We plan to overcome the latter approximation through direct studies of correlation functions of the energy momentum tensor of QCD on the lattice. So far this has only been done for SU(3) gauge theories [41]. We plan to perform calculations with light dynamical quarks on lattices of size $48^3 \times N_\tau$ with $N_\tau = 12,\ 16$ at several values of the temperature in the transition region, *i.e* for $T \sim (1-1.5)T_c$. As chiral symmetry is not a major concern for these studies, they can be performed within a staggered quark discretization scheme, and can partly utilize configurations generated for the study of the equation of state on $N_\tau = 12$ lattices. We estimate the cost for this project to be 100 Tflop-Years.

2.2.5 The phase diagram of strongly interacting matter

Lattice calculations currently provide the only *ab initio*, quantitative method for determining the phase diagram of strongly interacting matter. Figure 6 illustrates our present understanding at zero baryon number density (vanishing quark chemical potential). The extension of this diagram to nonzero baryon density is largely unexplored. Whereas at zero density the transition from the confined regime to the plasma regime is most likely a smooth crossover, it may happen that at nonzero density we encounter a genuine phase transition with dramatic consequences for heavy ion collisions and the structure of dense stars. Confirming the existence of a second order phase transition point at nonzero density and subsequently determining its location accurately can only be achieved through demanding numerical calculations. Experiments at RHIC and FAIR are under consideration that would search for this critical point. Quantitative predictions from lattice calculations are needed.

The question we will address is whether the critical line at unphysical quark masses in Fig. 6 moves outward toward physical quark masses as the chemical potential is increased. This question was addressed by de Forcrand and Philipsen [42] by introducing an imaginary chemical potential. We propose, instead, to reach nonzero chemical potential using a Taylor expansion. The calcu-

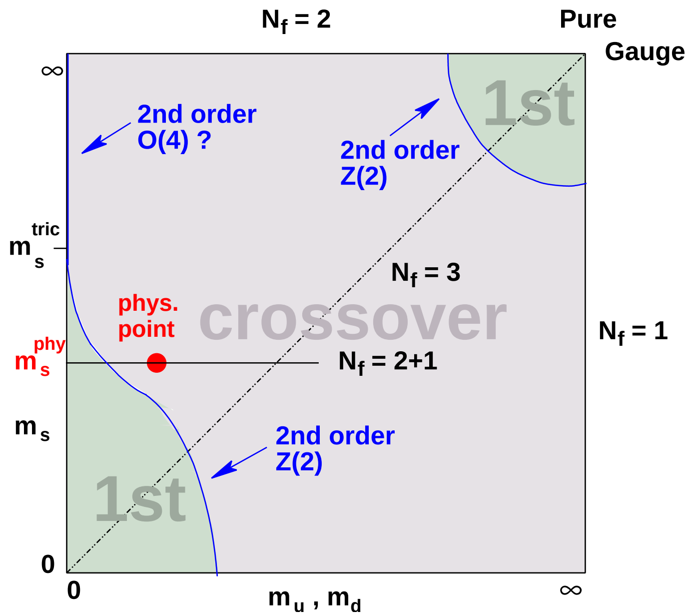

Figure 6: Sketch of the expected phase boundaries at zero chemical potential as a function of degenerate up and down quark masses. The physical point is plotted as a dot in the crossover region. To the left and below the 2nd order $Z(2)$ boundary, a high temperature first order phase transition occurs.

lation must be done at several carefully selected light quark masses. We estimate the cost of the calculation to be 100 Tflop-Years.

2.2.6 Conclusion

The advent of petaflops-scale computing promises dramatic gains in our understanding of the properties of strongly interacting matter at high temperatures and densities. We have described a program of lattice calculations that will (1) allow us to determine the equation of state of strongly interacting matter to an accuracy of 5%, (2) advance our understanding of the structure of the quark-gluon plasma, (3) determine some key transport coefficients, and (4) locate the critical surface of the QCD phase diagram at zero baryon density and predict its curvature as the baryon density is increased. The first goal will provide essential, solid input for hydrodynamical modeling of heavy ion collisions; the second will contribute to our understanding of the survivability of hadrons at high temperature; the third and most ambitious goal could very well give us the first reliable lattice result for a transport coefficient of the quark-gluon plasma; and the fourth could represent a potential breakthrough by moving us from a qualitative to a quantitative understanding of the phase diagram.

We have described a series of projects that can be tackled with resources of approximately 500 Tflop-Years. We place highest priority on the equation of state and the determination of transport coefficients.

Project	Lattice	T values	Quark Masses	MC traj.	TF-Yrs
EoS: $\mu = 0$, $T < 2T_c$	$48^3 \times 12$	10	2	100,000	110
EoS: $\mu = 0$, $T < 2T_c$	48^4	10	2	25,000	
EoS DWF: $\mu = 0$, $T \approx T_c$	$48^3 \times 10 \times 96$	4	1	50,000	80
EoS: $\mu > 0$ $T < 0.95T_c$ 8th order	$32^3 \times 8$	3-4	1	50,000	80
phase boundary $\mu \geq 0$	$32^3 \times 6$	4	5	10,000	100
spectral function, quenched	$128^3 \times N_\tau$	7	1	10,000	20
transport, dynamical	$48^3 \times N_\tau$	7	1	10,000	100

Table 5: CPU requirements in Tflop-Years for the thermodyanmics projects discussed in the text. Cost estimates are based on current experience at $N_\tau = 6$ and 8. The computational effort is assumed to scale with decreasing lattice spacing as a^{-11} with quark masses fixed in physical units. Simulations labeled $\mu = 0$ are at zero quark number density. Simulations labeled $\mu > 0$ imply a Taylor expansion in chemical potential to reach small nonzero densities. Temperature ranges are expressed in terms of T_c, the relevant crossover temperature. The parameter "MC traj." measures the size of the statistical sample needed. The lattice dimension for the domain wall fermion simulations (DWF) includes the "fifth dimension" L_s parameter.

2.3 The Spectrum, Structure, and Interactions of Hadrons

Understanding how the structure, spectroscopy, and interactions of hadrons emerge from QCD is one of the central challenges of contemporary nuclear physics. With recent investment in USQCD computer resources, advances in lattice field theory, and algorithmic developments, lattice QCD has entered a new era in which *ab initio* calculations of hadron structure observables can be compared directly with experiment. Hence, lattice QCD has now become an essential tool for nuclear physics, and with the necessary resources, it is poised to have major impact on contemporary experiments and on our fundamental understanding of hadron structure.

One central goal in hadron structure is precision calculation of fundamental experimental quantities characterizing the nucleon. These include form factors specifying the distribution of charge and current and how constituents interact to recoil together at high momentum transfer, moments of parton densities, helicity, and transversity distributions as a function of momentum fraction, and the moments of generalized parton distributions (GPD's). The calculations of all of these are specified as DOE Nuclear Physics 2014 milestones in Hadronic Physics (HP) [43]. These first principles calculations will have direct impact on key experimental HP milestones at TJNAF and RHIC-spin, including measuring the spin carried by glue in the proton and extracting accurate information on generalized parton distributions, determining EM and electroweak form factors and measuring flavor-identified q and $\bar{q}$ contributions to the proton spin. Another goal is obtaining insight into how QCD works: How does the spin of the nucleon arise from the helicity and orbital angular momentum of quarks and of gluons? What is the fundamental mechanism for confinement and what is the role of diquarks and instantons in hadrons? How does hadron structure change as one varies parameters that cannot be varied experimentally, such as the number of colors, the number of flavors, the quark mass, or the gauge group? Finally, one can exploit lattice QCD to provide information complementary to experiment. For example, the moments of GPD's calculated on the lattice and the convolutions of GPD's measured experimentally can be combined to constrain GPD's far more effectively than either experiment or lattice QCD could separately.

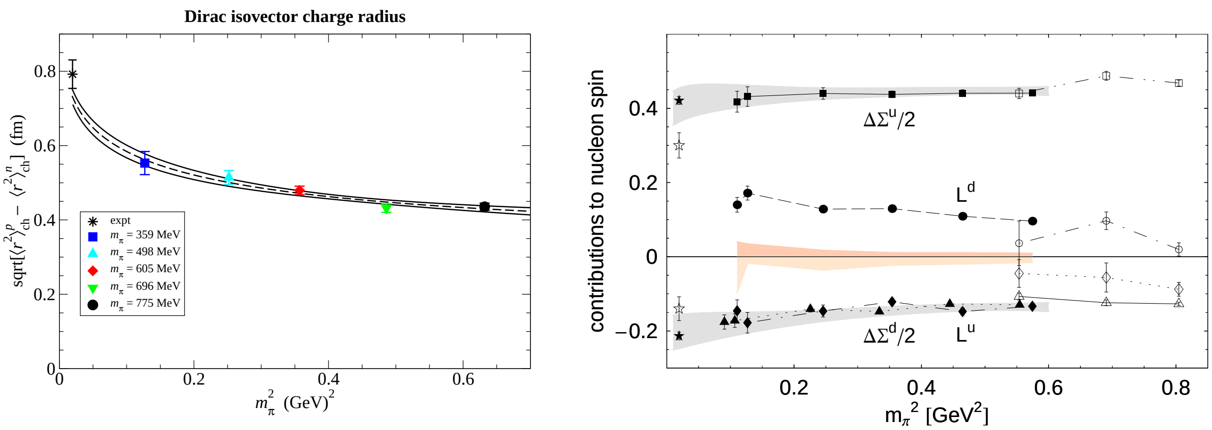

Figure 7: The left panel shows chiral extrapolation of the nucleon isovector charge radius, compared with experiment. The right panel shows the connected diagram contributions of the spin of the up and down quarks, $\Delta\Sigma^{u,d}$, and of the orbital angular momentum of the up and down quarks, $L^{u,d}$, to the total spin of the nucleon. The spin contributions are consistent with recent HERMES measurements (solid stars) and although the total orbital contribution is small, the u and d contributions are large and interestingly different from simple quark models.

A detailed knowledge of the meson and baryon spectra from first principles will distill the key degrees of freedom needed to describe the bound states of the theory. The complete combined analysis of available experimental data on the photoproduction of nucleon resonances and the measurement of the electromagnetic properties of the low-lying baryons are both HP milestones. The so-called hybrid mesons, a new form of excited state in which excitations of the gluon field play an explicit structural role, would be especially interesting, and the GlueX Collaboration proposal to seek information about exotic mesons is a flagship component of the 12 GeV upgrade at TJNAF. These intense experimental efforts have spurred a commensurate effort to predict and understand the hadron spectrum from first principles using lattice QCD.

A grand challenge for strong interaction physics is to be able to rigorously compute the properties and interactions of nuclei. The many decades of theoretical and experimental investigations in nuclear physics have, in many instances, provided a very precise phenomenology of the strong interactions in the non-perturbative regime. However, at this point we have little understanding of much of this phenomenology in terms of QCD. Acquiring this understanding is at the core of the HP milestone relating to a microscopic understanding of light nuclei, and the amalgam of lattice QCD and effective theories will be essential to its achievement.

2.3.1 Accomplishments

During the past five years, USQCD theorists have calculated a broad range of nucleon observables in full QCD for a range of quark masses down to the chiral regime, where the pion mass is as low as 350 MeV and chiral perturbation theory enables extrapolation to the physical pion mass. The first calculations with Wilson fermions were restricted to pions above 500 MeV and introduced the methodology for calculating the lowest three moments of quark, spin, and transversity distributions [44], electromagnetic and generalized form factors corresponding to the lowest three moments of generalized parton distributions [45, 46], and transition form factors between the nucleon and Delta to explore the role of deformation [47, 48].

An important step toward the chiral regime with light quarks was taken by introducing a hybrid action combining computationally economical staggered sea quark configurations generated by the MILC collaboration and domain wall valence quarks that have lattice chiral symmetry. Nucleon observables were calculated at five pion masses down to 350 MeV with a lattice spacing of a = 0.125 fm. Volumes of $(2.5\ \text{fm})^3$ and $(3.5\ \text{fm})^3$ were used at the lightest mass, and observables were extrapolated to the chiral limit using chiral perturbation theory [49, 50, 51, 52].

One salient accomplishment was calculating the nucleon axial charge, governing neutron β decay, to a precision of 6.8% and in agreement with experiment [49]. Another significant result was calculation of the isovector electromagnetic form factor, and as shown in Fig. 7, the isovector charge radius calculated from the slope of the F_1 form factor extrapolates via chiral perturbation theory close to the experimental value [53]. These hybrid action calculations also provide a first glimpse of the origin of the nucleon spin in the chiral regime. The contribution of the spin of an up or down quark q to the total spin of the nucleon is given by the zeroth moment of the spin dependent structure function $\frac{1}{2}\Delta\Sigma^q = \frac{1}{2}\langle 1\rangle_{\Delta q}$ and by the Ji sum rule [54], the contribution of the quark orbital angular momentum is given by $L^q = \frac{1}{2}\left(A^q_{20}(0) + B^q_{20}(0)\right) - \frac{1}{2}\Delta\Sigma^q$. The dominant connected diagram contributions to these matrix elements [52] are shown in the right panel of Fig. 7, where one observes that the chiral extrapolation of the spin contributions extrapolate to the experimental HERMES results and the orbital contributions are separately substantial and of the opposite sign from that given by the Dirac equation in a central potential. Additional accomplishments include calculation of the vector and axial transition form factors for the nucleon to Delta transition [55, 56, 57], yielding non-zero electric and Coulomb form factors revealing deformation, and use of partially quenched calculations to determine the neutron-proton mass difference [58].

Recently, more computationally expensive dynamical domain wall hadron structure calculations are also entering the chiral regime. The axial charge and first moments of the quark and spin distributions have been calculated with two flavors for pion massed down to 490 MeV [59, 60], and first results for vector and axial form factors, the tensor charge, and the first moments of the quark and spin distributions have been reported for 2+1 flavors at four masses from 330 to 670 MeV [61].

In spectroscopy, correlation matrix techniques offer a powerful means for determining the excitations of the theory, successfully demonstrated in the determination of the pure-gauge Yang-Mills glueball spectrum [62, 63]. The technique requires a basis of operators respecting the symmetries of the lattice. Such a basis for baryon states has been developed [64, 65], and their efficacy demonstrated in the quenched approximation to QCD [66, 67], illustrated in Fig. 8. In the meson sector, important progress has been made at identifying the spins of states in the approach to the continuum limit, and applying this method to resonance spectrum in charmonium [68]. The first calculation of the transition form factors between the lowest-lying charmonium states was performed, showing good agreement both with QCD-inspired models, and with experimental measurements [69], and the two-photon decay rate was computed [70], laying the ground for future studies of mesons composed of light (u/d) and strange quarks.

The foundation of the successes at studying the nature of the hadron-hadron interaction was the realization that lattice computations in the foreseeable future could be used to derive rigorous results for nuclear physics, opening up a new avenue for lattice computations that are at the core of this proposal [71].

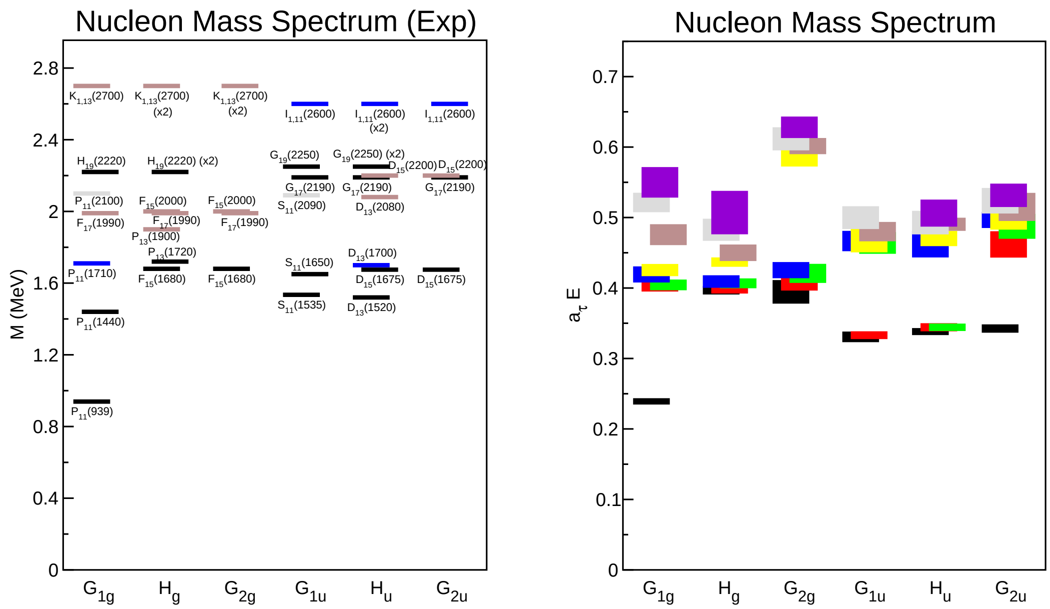

Figure 8: The left-hand panel shows the currently known spectrum from experiment, plotted according to the cubic symmetries allowed by the lattice; black denotes four-star states, blue denotes three-star states, tan denotes two-star states, and gray denotes a one-star state. The right-hand panel shows the nucleon spectrum in quenched QCD with $m_\pi \sim 700$ MeV from Ref. [66], where even at this unphysically large pion mass, one observes the emergence of patterns seen in the experimental spectrum.

Meson-meson scattering lengths were calculated with domain-wall fermions on $N_f = 2+1$ dynamical MILC sea configurations with pion masses down to $m_\pi \sim 290$ MeV [72, 73], obtaining the results shown in Figure 9. The first prediction of the $K\pi$ scattering lengths in both isospin channels was made possible by combining the lattice QCD calculation in the $I = 3/2$ channel with chiral perturbation theory, the result of which is also shown in Figure 9.

The first studies of nucleon-nucleon scattering with fully-dynamical lattice QCD were performed [74]. Although the pion masses employed were too large to uniquely match to low-energy effective theory, the calculation demonstrated the power of the method, which we will exploit in this proposal. In addition, the first calculations of hyperon-nucleon interactions were performed [75], in which it was shown that the scattering phase-shifts for elastic processes, such as $n\Sigma^-$, of importance for the nuclear equation of state at high densities can be extracted.

2.3.2 Future opportunities

Using the methodology and algorithms already developed and tested in calculations entering the chiral regime, the proposed new computational resources will enable precision calculation of wide range of isovector operators that can be evaluated with connected diagrams, including electroweak form factors, moments of structure functions, generalized form factors corresponding to moments of generalized parton distributions, and transition form factors, ultimately reducing the combination of systematic and statistical errors to the level of a few percent. Future calculations will utilize dynamical domain wall fermions, which will reduce systematic error and enable calculations far closer to the chiral limit. Between now and 2010, configurations will be generated with lattice spacing 0.086 fm at pion masses 330, 276, and 194 MeV and lattice volumes (2.7 fm)^3 and (4.1

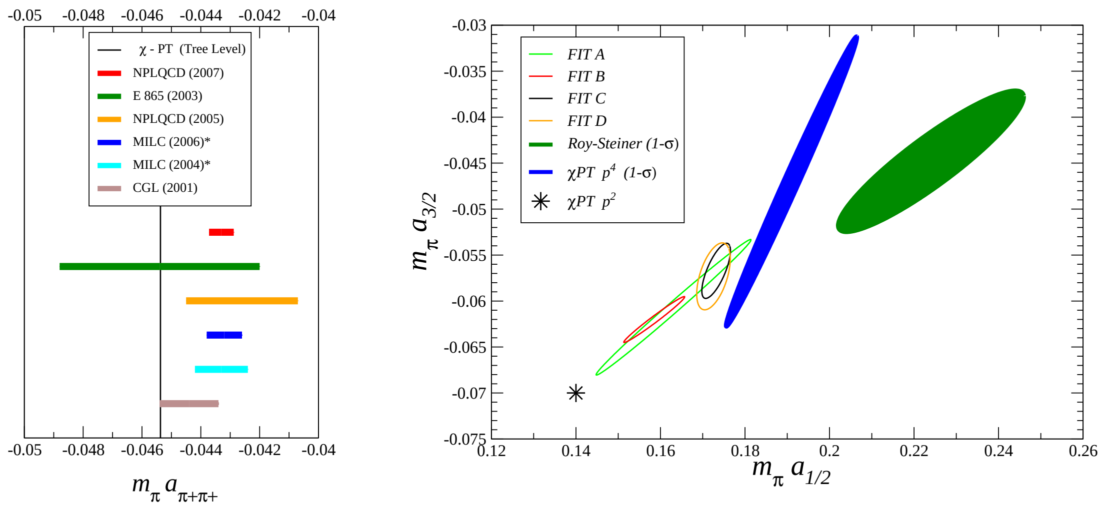

Figure 9: Meson-meson scattering lengths. The left-hand panel shows various determinations of the $I = 2\pi\pi$ scattering lengths; the red bar denotes the lattice computation in Ref. [73]. The right panel shows the prediction for the $K\pi$ scattering lengths at the physical pion mass [76].

fm)^3. As shown in Table 4, calculations from 2010 to 2014 will focus on two objectives. The first is to complete calculations with the same action at four masses with the high statistics required for precision calculations. The second is to utilize algorithmic improvements in a new action that will enable calculations down to the physical pion mass with more accurate enforcement of chiral symmetry and larger volumes. Note that for baryon observables, the amount of computer time required to analyze a configuration is of the order of half the time required to generate it.

In addition, the proposed resources will also enable crucial calculations on new, presently inaccessible observables. To calculate flavor-singlet matrix elements, in addition to the connected contributions, it is necessary to calculate disconnected contributions, which are typically several orders of magnitude more computationally expensive. Current development of eigenmode expansions and stochastic source techniques will enable calculation of flavor singlet form factors, moments of quark distributions, and generalized form factors, thereby addressing the full range of experimental nucleon observables including strangeness contributions. Building on the successful use of an improved gluon operator to calculate the gluon contribution to the pion momentum [77], the contribution of gluons to the nucleon mass, momentum and angular momentum, along with the associated mixing of gluon and flavor singlet operators, will be calculated. Additional opportunities that may be accessible include operator mixing of higher moments of structure functions and generalized form factors, higher twist operators, the neutron electric dipole moment, and differences between moments of structure functions of a free neutron plus a free proton and a deuteron.

In spectroscopy, the resources requested in this proposal will enable USQCD to capitalize on the achievements cited above to perform a program of computations that will provide insight into the resonance structure of QCD, and be crucial for achieving the DOE's investment in experimental resources.

The anisotropic Clover lattices tabulated in Table 6 are designed to enable high-precision investigations of the resonance spectrum. Lattices are generated at two volumes to enable the single- and multi-hadron states to be delineated, at two lattice spacings to enable the approach to the contin-

a (fm)	m_ℓ/m_s	Size	m_π (MeV)	L (fm)	$L\,m_\pi$	MC traj.	TF-Yrs
0.10	0.096	$48^3 \times 128$	220	4.8	5.3	10000	5.0
0.10	0.065	$48^3 \times 128$	180	4.8	4.3	10000	5.6
0.10	0.065	$64^3 \times 128$	180	6.4	5.8	10000	16.5
0.10	0.036	$64^3 \times 128$	135	6.4	4.3	10000	20.5
0.10	0.036	$80^3 \times 128$	135	8.0	5.4	10000	47.3
0.10	TOTAL						94.9
0.08	0.096	$48^3 \times 128$	220	3.8	4.2	10000	6.7
0.08	0.096	$56^3 \times 128$	220	4.5	4.9	10000	11.9
0.08	0.065	$56^3 \times 128$	180	4.5	4.0	10000	13.5
0.08	0.065	$72^3 \times 128$	180	5.8	5.2	10000	34.6
0.08	TOTAL						66.7

Table 6: CPU requirements in Tflop-Years (TF-Yrs) [26] for possible future configurations with anisotropic Clover quarks. "MC traj." gives the lengths of the runs in molecular dynamics trajectories. These lengths are chosen so that statistical errors should be sub-dominant for quantities of interest. All estimates are for two degenerate light quarks of mass m_l and a strange quark at its physical mass.

uum limit to be explored, and at decreasing values of the pion mass. In addition to the resources devoted to gauge-configuration generation, we will exploit recent progress in applying stochastic methods for the computation of "all-to-all" propagators [78], enabling many correlation-function measurements to be performed on each computation, and the computation of correlation functions for multi-hadron states. We estimate the computational cost of these measurements to be similar to those for gauge generation with the relative cost decreasing at decreasing pion mass.

In the first stage of this work, we will investigate the spectrum on volumes to 4 fm, at a lattice spacing $a = 0.1$ fm and at pion masses down to 220 MeV. This will yield the first calculations in the light-quark regime of the spectrum of exotic meson masses, and the first predictions of the π_1 hybrid photocouplings to conventional mesons, providing vital input into the GlueX experimental program. This regime will see the emergence of resonances, lying above two-particle decay channels into which they can, in the continuum, decay. Thus, measuring decay widths on the lattice will be an essential component of the program, and can be be accomplished using the volume dependence of the energy [79]. The efficacy of this method has been demonstrated recently in the computation of the width of the ρ mass [80]. The computation of the low-lying baryon resonance spectrum, including their decay widths, will confront the experimental analysis to yield insights into the degrees of freedom of QCD.

By the end of the third year of the project, ensembles will have been generated enabling computations of the spectrum at two lattice spacings, at two volumes and at pion masses down to 180 MeV. Thus the spectrum of baryons and mesons in the continuum limit will be known, and the transition form factors for many of these states will become accessible. At the culmination of the proposal, the first lattices at the physical value of the light-quark masses will be available, enabling the achievement of the goal of lattice computations of the spectrum for direct comparison with experiment.

The program to understand the hadron-hadron interaction imposes the same requirements as that to extract the spectrum and hadron structure, notably, the ability to resolve the spectral eigenvalues with high precision, the need for computations at several volumes to extract the scattering lengths, and the matching of the calculations to the chiral effective theory to describe interactions at the physical quark masses. The development of efficient stochastic solvers for the quark propagators will be especially beneficial, enabling multi-hadron states to be investigated.

Computations on the anisotropic lattices at volumes of 4 fm and pion masses down to 220 MeV will enable high-precision calculations of the $\pi\pi$, $K\pi$ and KK scattering amplitudes, and provide predictions that can confront the upcoming experimental determinations. Exploiting the lattices at pion masses down to 180 MeV will enable us to determine the nucleon-nucleon scattering lengths within the range of convergence of effective field theory with sufficient precision to enable an extrapolation to the physical pion mass.

Beyond the benchmark calculations discussed above, the hyperon-nucleon interaction is *terra incognita*. It has important astrophysical implications, in particular impacting the late time evolution of supernova. Thus, hypernuclear experimental programs are gaining increasing prominence. Furthermore, the hyperon-nucleon interactions probes an additional aspect of the nuclear force, the role of valence strange quarks in the force between hadrons. Thus, these resources will enable an important new avenue for lattice QCD to work in concert with a vigorous experimental hypernuclear program.

2.4 Lattice Gauge Theory for Physics Beyond the Standard Model

The first experiments at the Large Hadron Collider (LHC) will soon begin to probe physics at the TeV scale, attempting to unravel the dynamics of electroweak symmetry breaking and flavor physics. They are very likely to reveal new non-perturbative physics beyond the QCD sector of the Standard Model. Non-perturbative investigations of lattice field theory will play a crucial role in understanding the theoretical options and the experimental signatures. Theorists have proposed a wide variety of possible scenarios involving new gauge theories in the TeV region. These, for convenience, may be placed in three broad categories: Standard Model Higgs, SUSY (e.g. Minimal Supersymmetric Standard Model or MSSM), and new strong dynamics (e.g. new confining gauge theories such as technicolor). Some lattice field theory studies have been initiated in all these categories, but, as in the case of QCD, a serious effort requires substantial computational resources to carry out full scale simulations that include fermions in the chiral regime. Fortunately, the LHC era coincides with the arrival of multi-teraflop/s computers and highly improved algorithms. It is time to vigorously pursue lattice gauge theories for TeV physics beyond the Standard Model.

2.4.1 Higgs dynamics in the Standard Model

One scenario is the discovery of the Standard Model Higgs at the LHC in a fairly narrow mass range with little hint of its origin and without other new particles. This is often thought of as the least exciting option; however, this completion of the Standard Model spectrum raises serious issues in itself. First is the question of what is actually meant by the Standard Model in practice and its expected mass range.

A precise, if somewhat narrow, definition of the Standard Model only allows the well-known renormalizable operators in the Lagrangian. This description in perturbation theory ignores the intrinsic cutoff correctly included in the triviality scenario of the Higgs coupling [81, 82]. If the cutoff is at the Planck scale, the renormalizable Lagrangian is thought to be sufficient for all practical purposes, and the cutoff only comes into play in setting upper and lower bounds on the mass of the Higgs particle. The Higgs mass might fit into a narrow band in the 140-180 GeV range, hinting at a cutoff not far from the Planck scale. However, a Higgs mass significantly outside this range could signal the presence of nonrenormalizable operators in the effective Standard Model Higgs Lagrangian with low cutoff. For unexpected Higgs mass values new operators are needed because a low mass Higgs can trigger an instability due to its large Yukawa coupling to the top quark, whereas a large mass Higgs implies some new non-perturbative physics. Both features invoke the presence of higher dimensional nonrenormalizable operators on the TeV scale, but with constraints from electroweak precision data.

The Standard Model, notwithstanding all its successes, should then be viewed as an effective field theory with a cutoff much lower than the Planck scale, consistent with new physics in the TeV mass scale, perhaps within the reach of the LHC. To resolve these issues and to explore the options for non-perturbative effects in the TeV range it is essential to explore the Higgs-Top and Higgs-Top-QCD system in the TeV range non-perturbatively. Lattice Higgs simulations have been an important on going project for more than a decade [81, 82, 83] but only with the advent of Terascale computing and the application of chiral fermions can one explore coupled Higgs-Top dynamics properly [84].

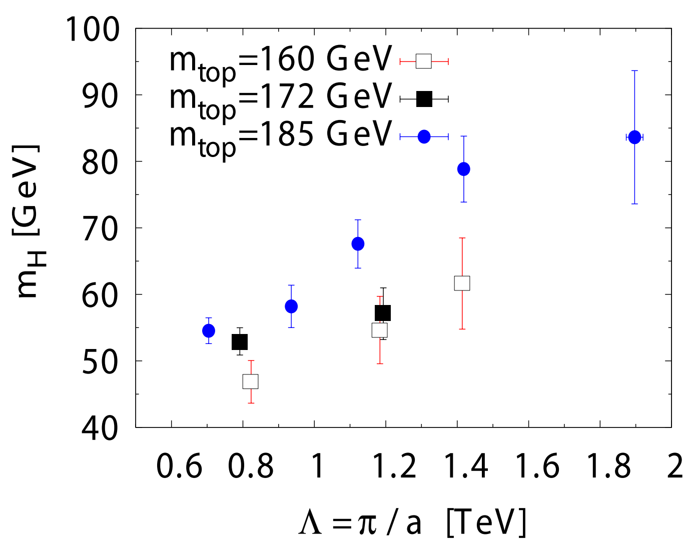

Figure 10: The lowest Higgs mass is plotted as a function of the lattice momentum cutoff for three different values of the Top mass. All simulation data are converted to physical units using $v = 246$ GeV for the vacuum expectation value of the Higgs field.

Some preliminary results are displayed in Fig. 10 for the Higgs mass lower bound in full dynamical Monte Carlo simulation of the Higgs-Top sector with overlap fermions in the limit of vanishing Higgs self-coupling at the cutoff scale [84]. It will be important to explore the Higgs-Top phase diagram for negative Higgs self-coupling on the cutoff scale while dimension six operators provide the stability.

A new generation of simulations will be required with increased resources to obtain phenomenologically important results with the lattice artifacts under full control. To stay close to the critical line of spontaneous symmetry breaking will require the vacuum expectation value of the Higgs field to be around 0.2 in lattice spacing units, which corresponds to a lattice momentum cutoff of approximately 3 TeV. To explore Higgs masses in the 100 GeV to 200 GeV range will drive down the Higgs mass to the 0.1 range, which requires spatial lattices with 40 or 50 links. To explore the heavy Higgs particle scenario will be similarly demanding. In addition, accurate Higgs physics requires a much larger number of lattice configurations than fermionic measurements in lattice QCD. The representative lattices in Table 7 are given for a Higgs mass sequence of $M_H = 0.15, 0.10, 0.08$ with the linear box size L about 5 times the Higgs correlation length.

2.4.2 Supersymmetric field theories

A second scenario involves the discovery of supersymmetry with its attendant zoo of new particles. A prime focus of initial work on lattice supersymmetric theories will be $\mathcal{N} = 1$ super Yang-Mills. It is perhaps the simplest example of a supersymmetric theory and constitutes an important part of the minimal supersymmetric extension to the Standard Model.

This theory is predicted to have a discrete Z_N chiral symmetry which has been suggested to break spontaneously to Z_2. First generation numerical simulations of the dynamical theory have already been performed using domain wall fermions [85]. The breaking of the $U(1)_R$ symmetry down to Z_2 is indeed observed. The presence of fractional topological charge on a toroidal lattice was also observed in the quenched theory using the overlap method. These studies demonstrate that the lattice can in fact explore the difficult non-perturbative questions regarding this theory.

The next steps for $\mathcal{N} = 1$ super Yang-Mills are to determine in detail the pattern of $U(1)_R$ symmetry breaking, give accurate masses of the low lying spectrum and study mixing of the pseudoscalar glueball and the eta prime meson. All of these issues have been studied theoretically on the basis of various conjectures and approximation schemes. For $\mathcal{N} = 1$ super Yang-Mills, there are no Nambu-Goldstone bosons; so the box size L should be larger that the lightest mass scale $L \gg O(1/m_V)$. The main effort is to make the gluino mass as light as possible while taking the continuum limit [85] and controlling finite volume effects. The cost shown in Table 7 is for the SU(3) gauge group. For other numbers of colors N_c, the cost should scale as $\left((N_c^2-1)/8\right)^{1.5}$. An improved quantitative understanding of this theory is the first step toward a better understanding of super-QCD and similar theories. If the LHC reveals evidence for new strong dynamics at the electroweak scale it is important to understand whether that dynamics corresponds to a supersymmetric gauge theory.

A much larger range of SUSY theories are beginning to be considered on the basis of elegant lattice constructions using ideas drawn from orbifolding in string theory and the twisting procedure used in constructing topological field theories [86, 87, 88]. These formulations retain a degree of exact supersymmetry and the hope is that this will dramatically reduce the amount of fine tuning required to obtain the correct continuum limit. They lead to surprising lattice geometries such as that illustrated in Fig. 11, where the fermionic partners are scattered on the lattice in manner reminiscent of staggered fermions, but with no unphysical degrees of freedom. Initial Monte Carlo simulations of these models have already yielded interesting non-perturbative information on the vacuum structure and can be used to probe issues of spontaneous supersymmetry breaking [89].

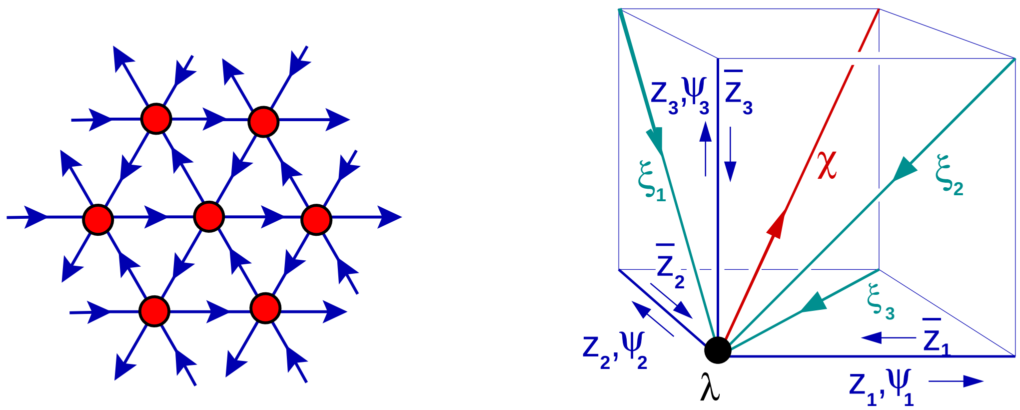

Figure 11: On the left is the lattice for supersymmetric Yang Mills in $d = 2$ with $\mathcal{Q} = 16$ supercharges. On the right is lattice for supersymmetric Yang Mills in $d = 3$ with $\mathcal{Q} = 8$ supercharges. The z_i are bosons, while the other fields are one-component fermions

These theories are also very interesting from a string theory point of view. They are thought to correspond to ten-dimensional string theories via a set of dualities that generalize the original AdS/CFT correspondence of Maldacena. This is an area where lattice simulations could play an important role in casting light on the non-perturbative phase structure of string theories. Initial work is already allowing us to probe black hole thermodynamics using simulations of thermal Yang-Mills systems [90]. We anticipate extending these lattice studies to the most interesting case of $\mathcal{N} = 4$ super Yang-Mills in four dimensions within the time frame of this proposal. Code for simulating this theory is already under development.

2.4.3 New strong dynamics

A third scenario is the discovery of a new strong dynamics for electroweak symmetry breaking, which could provide an ideal application for lattice field theory. In this case the Higgs phenomena may well be most cleanly described as a composite arising in a new strongly coupled gauge field theory, such as that proposed in technicolor, Higgsless models, and extra-dimensional (Randall-Sundrum) models. Other strong coupling methods, usually motivated by the AdS/CFT conjecture, can give only qualitative results, so lattice field theory offers the only *ab initio* non-perturbative method for making quantitative predictions.

In technicolor models, only two flavors are needed to provide the Goldstone modes for the massive electroweak longitudinal vectors via the Higgs mechanism. Technimeson resonances also occur in these models with experimental signatures that may mimic the Higgs boson of the Standard Model. To address the flavor problem and provide masses for Standard Model fermions, technicolor has to be "extended" with additional techniflavors and new interactions with Standard Model fermions. However, precision electroweak constraints severely limit the number of additional technifermion flavors unless their mass scale is pushed into the multi-TeV range, well above the typical electroweak scale of minimal technicolor. The conjectured mechanism that allows light electroweak bosons and additional multi-TeV fermions to coexist while evading electroweak constraints is for the additional heavy fermions to slow down the "running" gauge coupling, leading to so-called "walking" evolution in the infrared (IR). Indeed it has been known for a long time that sufficient numbers, N_f, of flavors leads to a Banks-Zaks conformal fixed point [91] in the IR, so that the Yang-Mills theory no longer confines or breaks chiral symmetry, leading to a so called the non-Abelian Coulomb phase or conformal window. The goal of a "walking" theory is to be just below the minimum number of flavors required to enter the conformal phase. To determine the minimum number and to understand the dynamics below this point is highly non-perturbative problem. Recently, a lattice study [92] using the Schrödinger functional approach for 4-plets of staggered fermions, indicates that the transition to conformal behavior for SU(3) Yang-Mills requires between 8 and 12 fundamental massless flavors, as illustrated in Fig. 12.

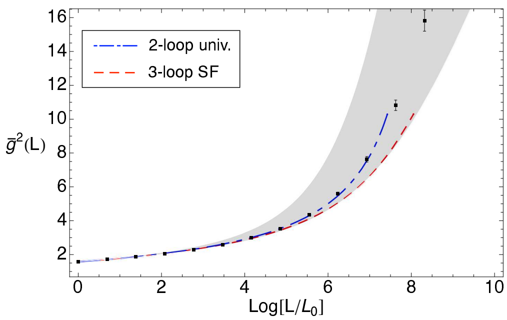

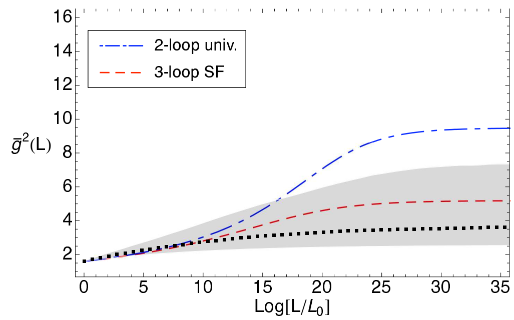

Figure 12: Comparison for the running coupling from step scaling for $N_f = 8$ on the left and $N_f = 12$ on the right indicating the critical number of flavors for the formation of the conformal window lies in this interval.

The next step in this project is to accurately determine the edge of the conformal window with domain wall fermions and to perform lattice calculations below the critical number of flavors to search for signals of walking dynamics. For example in Table 7 a series of lattices are proposed for $N_f = 8, 9, 10, 11$ holding fixed the lightest non-Nambu-Goldstone boson mass $am_V \ll 1$. It is also possible that "walking" may not occur in SU(3) Yang-Mills with fermions in the fundamental representation or that the required number of fermions is not appropriate phenomenologically. Since "walking" may occur in Yang-Mills theories with fermions in higher representations of the gauge group, it is natural to explore these as well. There is some evidence [93] that near conformal behavior might occur already for $N_f = 2$ flavors for fermions in the **6** or **8** representations of SU(3).

In studying theories close to a conformal fixed point one must be aware of the danger of enhanced finite volume effects. This calls for a series of lattices of various volumes and numbers of flavors approaching the conformal edge to study and even to exploit the finite size effects in measuring physical parameters. The goal of this project is to understand flavor dependence of the spectra, particularly the vector ρ and axial-vector a_1 mesons, the chiral condensate and the analog of f_π. The spectrum of vector and axial-vector mesons are essential ingredients for the study of precision electroweak constraints. These studies will reveal if it is possible to avoid conflict with precision electroweak experiments and still be able to build models for LHC events.

Finally the nature of the conformal theory associated with an IR fixed point [91] appears interesting in its own right. While Seiberg duality provides considerable detail concerning the conformal phase in supersymmetric QCD, essentially nothing is know about strongly coupled conformal theories without supersymmetry. Indeed, as noted recently [94], the experimental possibility of such a conformal sector has not been ruled out.

2.4.4 Computational resources

Here we give some examples of the lattices needed to pursue this research. The specific choice of lattices for these relatively unexplored physics topics will change in response to experience gained year by year and with the discoveries at the LHC. In general all three scenarios described above will rely on chiral (Domain Wall or overlap) fermions with a variety of fermionic degrees of freedom, and extrapolations to the chiral limit. Estimates of these computational requirements are based on experience with similar 2+1 flavor QCD lattices and benchmarks for both multi-flavor QCD and $\mathcal{N} = 1$ super Yang-Mills, using codes built on the SciDAC API for message passing. The cost scales roughly as $N_f^{1.5}$ and as $d(R)^{1.5}$ where $d(R)$ is the dimension of the representation. To simplify Table 7, we have only listed the largest, most computationally expensive lattices. However, due to the exploratory nature of this research, it is expected that this project will make good use of smaller lattices in its early stages. We estimate that these smaller lattices will account for an additional cost of less than 50%.

N_f	am_ℓ	Size	am_V	Lm_V	"MC traj."	TF-Yrs
HIGGS	0.001	$32^3 \times 32 \times 32$	0.167	4	10000	0.66
SUSY	0.001	$24^3 \times 24 \times 32$	0.167	4	5000	0.66
$N_f = 8,9,10,11$	0.0075	$24^3 \times 64 \times 24$	0.25	6	4×5000	2.32
TOTAL						3.64
HIGGS	0.00075	$48^3 \times 48 \times 32$	0.125	4	10000	2.11
SUSY	0.00075	$32^3 \times 32 \times 32$	0.125	4	5000	2.11
$N_f = 8,9,10,11$	0.005	$32^3 \times 64 \times 24$ u	0.25	8	4×5000	6.80
TOTAL						11.0
HIGGS	0.0005	$64^3 \times 64 \times 32$	0.125	6	10000	23.7
SUSY	0.0005	$48^3 \times 48 \times 32$	0.125	6	5000	23.7
$N_f = 8,9,10,11$	0.0035	$48^3 \times 96 \times 32$	0.25	8	4×5000	66.0
TOTAL						113.4

Table 7: Representative lattice ensembles for the Top-Higgs dynamics (with one fundamental fermion), for $\mathcal{N} = 1$ SUSY (with one adjoint fermion) and for SU(3) Strong dynamics (with $N_f = 8,9,10,11$ fundamental fermions).

Finally we emphasize that, in contrast with QCD, the lack of experimental data for these theories makes distinguishing lattice artifacts and continuum predictions more challenging and substantially raises the standards for obtaining convincing results. For this reason, we have recommended beginning with models that are for the most part close variants of QCD, enabling comparison with QCD to bring confidence to the methodology. Also, this makes sense from a model building perspective as a way to test the "conventional wisdom" regarding the behavior of non-perturbative QCD-like theories as the gauge or matter content is varied, which is often based on simple scaling assumptions. Validating or modifying these non-perturbative heuristics will have a substantial impact on comparing these models with LHC data.

However these conservative priorities could change dramatically in the future. For example new algorithms are being developed for operators that couple to "disconnected" or vacuum diagrams [95] that, for example, allow the strangeness content of the nucleon to be measured, which is the major uncertainty in determining the possibilities for detection of the dark matter candidate in the

MSSM[96] at the LHC. Disconnected diagrams are even more important in studying the SUSY spectrum. Other examples of theoretical challenges include (1) the "sign problem" encountered not only at finite chemical potential but also in extended supersymmetry, (2) spectral studies of unstable particles, be they the rho or the Higgs and (3) the representation of chiral gauge theories on the lattice [97], which is perhaps the most important problem for understanding a non-perturbative treatment of the Standard Model and beyond. Work is in progress to address all of these challenges. The exploratory use of lattice techniques for quantum field theory may yield new unexpected surprises and insights, potentially leading to the development of powerful new methods, along with a deeper insight into the special properties of non-perturbative field theory.

3 Hardware Requirements

To reach the scientific objectives set out in Section 2 will require both access to the DOE's leadership class computers and the acquisition of computers dedicated to the study of QCD. The purpose of this proposal is to obtain funds to acquire and operate dedicated machines. We have applied separately for access to the DOE's leadership class computers, the Cray XT4 and its successors at Oak Ridge National Laboratory (ORNL), and the IBM BlueGene/P at Argonne National Laboratory (ANL), through the DOE's Incite Program. Because there is close coupling between the work to be done on the two types of computers, we describe our plans for both here.

Lattice QCD calculations proceed in two steps. In the first one performs Monte Carlo calculations to generate gauge configurations with a probability proportional to their weight in the Feynman path integrals that define QCD. These configurations are stored, and in the second step they are used to calculate a wide variety of physical quantities. The same configurations are often used to study problems in high energy physics and in nuclear physics. Configuration generation is the most computationally intensive part of our work, and in most cases limits the rate of progress. Since it involves a Markov chain, it must be carried out in a small number of streams. The generation of gauge configurations with small enough lattice spacings and quark masses to reach the levels of accuracy discussed in Section 2 requires computers that enable one to apply very large numbers of processors to individual calculations. The same is true of the calculation of propagators for light mass quarks. These two types of calculations are therefore best done on leadership class computers. On the other hand, dedicated computers, such as the clusters acquired during LQCD, are the appropriate platforms for those aspects of our work that require large computational resources, but do not require that the full power of a leadership class computer be applied to individual jobs. Much, but not all, of the physics analysis performed on stored gauge configurations falls into this category. Although the total number of floating point operations used in the analysis of a given ensemble of gauge configurations ordinarily equals or exceeds the number needed to generate the ensemble, individual analysis jobs are typically smaller, and can be run efficiently on fewer processors than configuration generation jobs. The analyses of different configurations are independent of each other, so they can be run in parallel. Other aspects of our work that are best suited for dedicated computers are the study of high temperature QCD, the generation of less challenging gauge configurations, the development of new algorithms, new formulations of QCD on the lattice, and other exploratory studies. Such work plays a very important role in our research, requires large computational resources, but does not require the capabilities of leadership class machines.

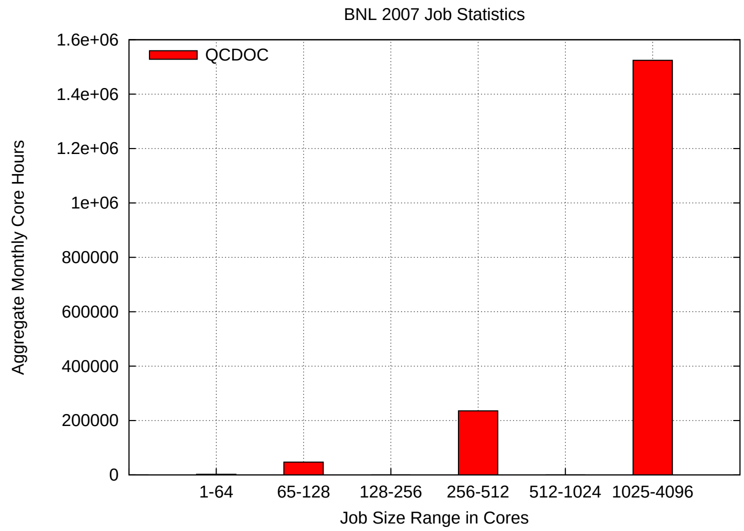

Figure 13: The aggregate core hours used per month on the QCDOC computer at BNL as a function of the number of processor cores employed in the jobs. Core hours are normalized to those of 6n, an Infiniband cluster with a sustained performance of 1.07 Gflop/s per core.

The history of our field indicates that at least as much is gained by advances in algorithms as by advances in hardware technology, and we expect this trend to continue. Indeed, in the last few years the introduction of the Rational Hybrid Monte Carlo (RHMC) algorithm [98] has reduced the number of floating point operations needed to generate gauge configurations with light quarks by factors of four to eight compared to algorithms in use a few years ago. The gauge configurations we expect to generate during the course of LQCD-ext could not be produced by the proposed resources in any reasonable amount of time without the RHMC algorithm. We are confident that our continued work on algorithms will significantly enhance our productivity, extending the range of science we will be able to do. For example, members of USQCD are currently investigating a variety of approaches to accelerating the calculation of quark propagators, which have the potential to greatly enhance both configuration generation and physics analysis. Work of this type requires significant computational resources, but is not appropriate for leadership class machines.

We believe that we will do the most science by using the computers that are best suited for each phase of our work. Moreover, we cannot simply transfer all of our calculations to the leadership class machines because the 2008 Incite Call for Proposals specifically states that "*Applicants must also present evidence that they can make effective use of a major fraction of the processors of the high performance computing systems offered for allocation.*" This is certainly the case for configuration generation and the calculation of light quark propagators, but it is not true for much of our physics analysis or our exploratory studies. A few years ago, the bulk of the floating point operations in any lattice QCD calculation went into the generation of gauge configurations, but that is no longer so. In the 2007 USQCD allocation process, four projects that received allocations were deemed appropriate for leadership class computers. They received approximately 40% of the available resources. On the other hand, twenty-five projects, which received 60% of the available resources, were considered appropriate for dedicated computers. The latter projects were critical for achieving our scientific goals.

Figure 13 shows the core hours used per month on the QCDOC computer at BNL, as a function

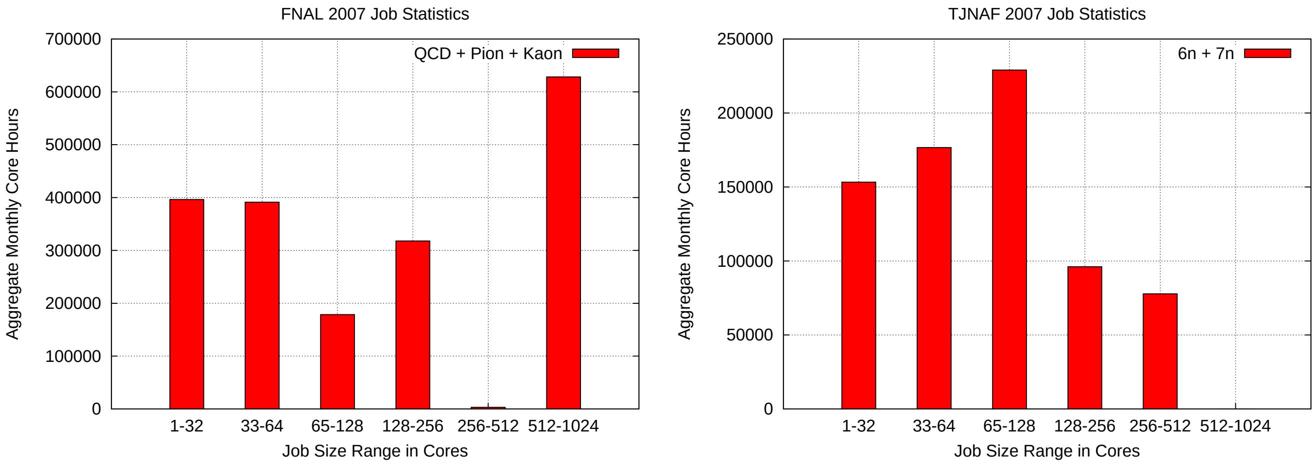

Figure 14: The aggregate core hours used per month on the clusters at FNAL (left panel) and at TJNAF (right panel), as a function of the number of processor cores employed in the jobs. Core hours are normalized to those of 6n, an Infiniband cluster with a sustained performance of 1.07 Gflop/s per core.

of the number of processor cores employed. Figure 14 shows the same quantity for the clusters at FNAL and TJNAF. These results are monthly averages for 2007. The core hours are normalized to those of 6n, an Infiniband cluster acquired by TJNAF in 2006. It has a sustained performance of 1.07 Gflop/s per core. Here and throughout this proposal, performance is measured as the average of that sustained by the sparse matrix inversion routines for computing the quark propagators for the Domain Wall and Improved Staggered (Asqtad) quark actions under production conditions. These routines consume a significant fraction of the floating point operations in our calculations, and are representative of the overall performance of our codes. The bulk of the jobs on the QCDOC were for configuration generation or the calculation of quark propagators, and the 1024 core jobs at FNAL were for configuration generation. We expect that by 2010 such jobs will be transferred to leadership class machines. The remaining jobs on the FNAL and TJNAF clusters are for physics analysis, as well as for algorithm development and other exploratory projects. They will certainly grow in size over time, but are unlikely to reach a size that would warrant moving them to leadership class computers. This data is one indication of the importance of dedicated computers the moderate-sized jobs required for our physics program.

We are convinced that the trend towards using an ever greater proportion of computing resources on physics analysis will continue. There are several reasons for this shift. As the capabilities of leadership class computers increase, it will be possible to generate configurations with smaller lattice spacings and quark masses, using more sophisticated formulations of lattice QCD. More resources are needed to analyze such configurations than the ones currently being generated. In addition, more realistic gauge configurations, coupled with greater computing power for analysis, will enable new calculations that could not previously be undertaken. In drawing up the plans presented in this proposal we have projected that analysis will continue to require more computing resources than lattice generation.

The computing cycles we receive from the LQCD dedicated computers and from our current Incite grant are heavily leveraged. As is discussed in Section 4, we have long standing collaborations with colleagues in the United Kingdom in our zero temperature studies with DWF and improved staggered quarks, with colleagues in Ireland in our work on the hadron spectrum, and with colleagues in Germany in our work on high temperature QCD. At present, the computing cycles contributed to our joint work by these collaborators, coupled with those from the Japanese funded Riken Brookhaven Research Center (RBRC), the National Energy Research Scientific Computing Center (NERSC), the National Science Foundation (NSF) supercomputer centers, and, most recently, the National Nuclear Security Administration (NNSA) BlueGene/L at Lawrence Livermore Laboratory (LLNL), roughly equal those provided by LQCD and the Incite Program. In planning for the period 2010–2014 we have assumed that roughly half of our computing cycles will continue to come from these sources.

A total of 725 Tflop-Years [26] are needed to generate the zero temperature gauge configurations enumerated in Section 2, half of which we plan to request on leadership class computers. We plan to ask for an additional 160 Tflop-Years on these machines for the generation of light quark propagators, bringing our total request for leadership class computing resources during the five year period of the proposal to 500 Tflop-Years, as is indicated in Table 8. As stated earlier, our major analysis projects required approximately the same number of cycles as generation of the gauge configurations, so we request approximately 165 Tflop-Years on dedicated computers for this aspect of our work, in addition to the 160 Tflop-Years for generating quark propagators. The total resources required for the thermodynamics projects described in Section 2.2 is 490 Tflop-Years. We plan to use 165 Tflop-Years on dedicated computers for these projects, with the remainder coming from other resources. Finally, we have budgeted a total of 85 Tflop-Years for exploratory projects and algorithm development, bringing the total resources from dedicated computers to 415 Tflop-Years. In Table 8 we set out the yearly usage of dedicated hardware and leadership class computers over the five years of the proposal. The dedicated hardware can be acquired with a fixed budget of $2.01 million per year.

Fiscal Year	Dedicated Hardware (Tflop-Years)	Leadership Class Machines (Tflop-Years)
2010	25	30
2011	45	50
2012	65	80
2013	110	135
2014	170	225

Table 8: Computing resources from the use of dedicated hardware (column 2) and leadership class computers (column 3) needed to carry out the scientific program set out in Section 2 by fiscal year. Computing resources are given in sustained Tflop-Years [26, 27].

3.1 Leadership Class Computers

We believe that the time we anticipate on leadership class computers is realistic. We have demonstrated that our codes scale to thousands of processors on these machines. Two examples are shown in Fig. 15, where we plot the total performance of the conjugate gradient routines for Asqtad quarks

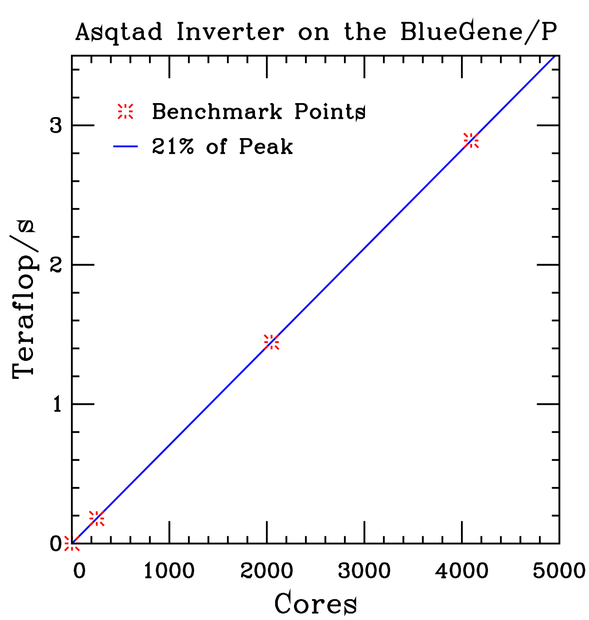

Figure 15: Performance in teraflop/s of the MILC Asqtad inverter on the BlueGene/P (left panel) and the Chroma anisotropic Clover inverter on the Cray XT4, as a function of the number of processor cores. The red bursts are the benchmark data points and the blue curves are straight lines from the origin through the data points with the largest number of cores.

on the BlueGene/P and for anisotropic Clover quarks on the Cray XT4 in teraflop/s, as a function of the number of cores. These are weak scaling tests in which the number of lattice sites per core is held fixed as the number of cores is increased. The red bursts are the benchmark data points, and the blue curves are straight lines drawn from the origin through the benchmark points with the largest number of cores. The percentage of peak that this line corresponds to is shown on the figure. We are confident that scaling persists for far larger numbers of cores than are shown in Fig. 15. Fig. 16 demonstrates near perfect weak scaling for the conjugate gradient routine of the Wilson quark action through 131,072 cores, the full size of the LLNL BlueGene/L on which this test was run. The software used in this test was a version of the Columbia Physics System (CPS) code optimized for the BlueGene/L by Pavlos Vranos [99] using native BlueGene communication calls, which are being integrated into our production codes for the BlueGene/P. It achieved 70.5 teraflop/s on the full machine, and led to a Gordon Bell Special Achievements Award.

Over the last few years members of USQCD have been allocated approximately 10% of the cycles on the NERSC IBM SP, Seaborg, and approximately 15% of the cycles in the first allocation of the NERSC Cray XT4, Franklin. In the current year our collaboration has an allocation through the INCITE Program of 10 million core hours on the ORNL Cray XT3/4, Jaguar, which amounts to approximately 10% of the available cycles, and we have submitted a proposal to the Incite Program seeking to obtain approximately 10% of the cycles on the leadership class computers at ANL and ORNL for 2008–2010. To see what this might mean in terms of scientific output, we note that the President's FY 2008 budget calls for the installation of a leadership machine at ORNL with a peak performance of 1,000 teraflop/s, and another at ANL with a peak performance of 250 to 500 teraflop/s by the end of FY 2008. Using the mid-point in the range of the ANL machine, there would be 1,375 teraflop/s (peak) available for allocation in FY 2009. Although we are aware that the Office of Science has more ambitious plans for expanding its leadership class facilities, we

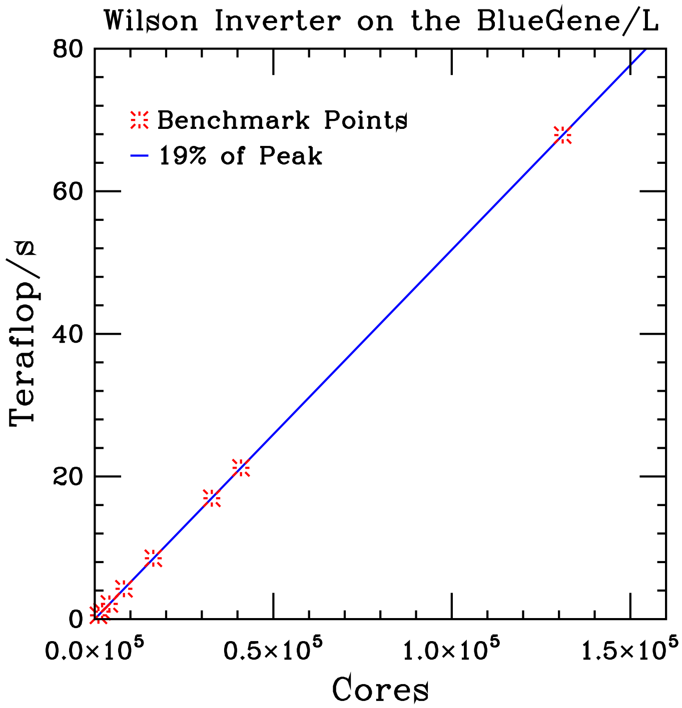

Figure 16: Performance in teraflop/s of the CPS Wilson inverter on the BlueGene/L as a function of the number of processor cores. The red bursts are the benchmark data points, and the blue curves are straight lines from the origin through the data points with 131,072 cores.

make the conservative assumption that they will grow in accordance with Moore's law between FY 2010–2014. Our full configuration generation codes currently sustains 15% to 20% of peak on these platforms. We believe that this performance will improve in the future, but use the lower figure in the following estimates. Under these assumption we would obtain a throughput from leadership class machines of 30 Tflop-Years in FY-2010 growing to 210 Tflop-Years in FY 2014. We believe that our assumptions regarding the leadership class machines are quite conservative. If they increase in performance more rapidly than we have projected, or if we obtain more than 10% of the time allocated on them, then our research will accelerate proportionately.

3.2 Dedicated Hardware

Our proposed process for the acquisition of dedicated hardware follows that of the current LQCD project. That is, each year we will acquire the hardware that best advances our science. As in LQCD, we propose to locate the hardware at BNL, FNAL and TJNAF. Our estimates of the price/performance of dedicated hardware is based on our experience with commodity clusters over the last seven years. From 2001 through 2005, we acquired a series of prototype clusters under our SciDAC I grant. In doing so we demonstrated that by carefully choosing components to optimize the performance of our codes, we could obtain highly cost effective computers for the study of lattice QCD. Three of these clusters, with a combined throughput of 1.65 teraflop/s are currently being operated under LQCD. During the first two years of LQCD we have acquired three commodity clusters with Infiniband networks: the 6n and 7n clusters located at TJNAF and the Kaon

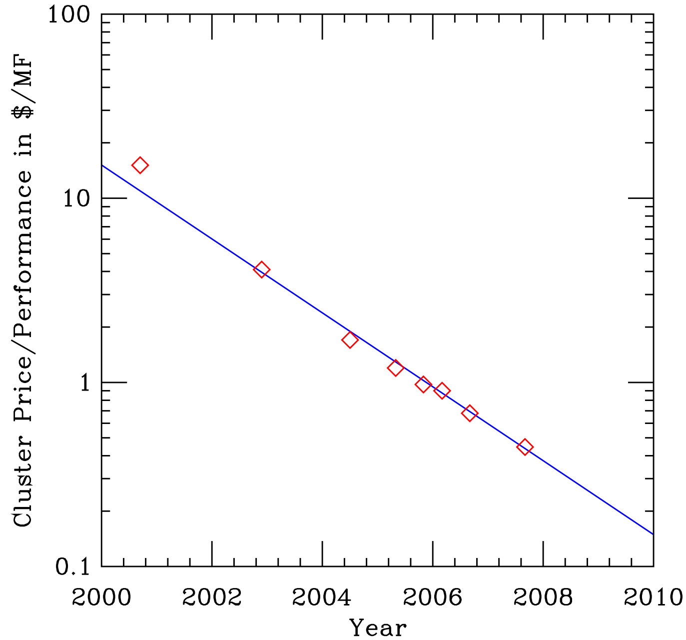

Figure 17: The price/performance in dollars per sustained Mflop/s of clusters built under the SciDAC I Program and the Lattice QCD Computing Project as a function of time. The red diamonds are the price/performance of individual clusters, and the blue line is a least squares fit to all but the earliest cluster point with the halving time of the price/performance fixed at 1.5 years.

cluster located at FNAL. These machines have a combined throughput of 6.14 teraflop/s. Since we began acquiring clusters under the SciDAC I Program, their price/performance has decreased in accordance with Moore's law with a halving time of 1.5 years, as is illustrated in Fig. 17, where we show the price/performance as a function of time on a log plot. The latest point on this plot is for the 7n cluster, and corresponds to a price/performance of $0.45 per sustained Mflop/s.

We have based our projections for dedicated computers on our experience with commodity clusters. We assume that Moore's law will continue to hold for clusters with gains coming primarily from the increase in the number of cores per processor, rather than the increase in clock frequency. If so, the price/performance will reach $0.15 per sustained Mflop/s in 2010 and $0.024 in 2014. Of course the extrapolation to 2014 has considerable uncertainty attached to it. We have found that the useful lifetime of a cluster is a minimum of three years, so in 2010 the clusters built under LQCD in 2007-2009 will still be in service. We also assume that the QCDOC will continue in operation through 2011. Table 9 shows the computers that will be carried over from LQCD, their location, throughput, and last expected year of operation. The JPsi cluster is to be built at FNAL with combined LQCD hardware funds for 2008 and 2009.

The throughput for dedicated hardware in the second column of Table 8 was obtained by adding that of the machines in Table 9 to the computers we expect to acquire under LQCD-ext with the proposed hardware budget of $3.0 million per year, given our projections regarding price/performance. We again assume that the lifetime of the new computers will be at least three years. Although these projections are based on our experience with clusters, we emphasize that in each procurement we will obtain whatever hardware best advances our science. Just as in the earlier LQCD project,

Cluster	Location	Throughput (Sustained Tflop/s)	Final Year
7n	TJNAF	3.0	2010
QCDOC	BNL	4.2	2011
JPsi	FNAL	6.2	2011

Table 9: Computers operated under LQCD that will be carried over to LQCD-ext. The first column gives the computer name, the second its location, the third its total throughput, as measured by our standard benchmarks, and the fourth the last year in which it will be operated.

we expect that the most efficient dedicated resources for lattice QCD will be a mixture of cluster hardware as discussed above and machines targeting favorable cost performance through a greater degree of custom-design. Such a machine, a possible future member of the BlueGene series, is presently being designed by a collaboration of physicists at Columbia and Edinburgh supported in part by the RBRC and working together with colleagues at the IBM corporation. We expect that the result of this effort will be a highly cost-effective computer that will run lattice QCD very efficiently. Given the direct participation of members of our collaboration in this effort and the long-term experience of BNL in supporting machines of this sort, we anticipate that a portion of the dedicated hardware funded through this proposal may be provided in this way. The performance figures in the second column of Table 8 therefore represent a lower bound on what we expect to achieve.

The proposed budget for LQCD II is given in Table 10. The hardware budget is fixed at $2.01 million per year, so the increases in throughput for the dedicated hardware come entirely from Moore's law. The operations costs shown in the third column of Table 10 are overwhelmingly salaries of laboratory staff, who will maintain the existing hardware and conduct the procurement of new hardware.

Fiscal Year	Hardware Budget	Operations Budget	Total Budget
2010	2.01	1.13	3.14
2011	2.01	1.33	3.34
2012	2.01	1.62	3.63
2013	2.01	1.37	3.38
2014	2.01	1.68	3.69
Totals	10.05	7.13	17.18

Table 10: Proposed budget for LQCD II in millions of dollars.

4 The USQCD Collaboration

Two important functions of the USQCD Collaboration are to set scientific goals for research in lattice QCD in the United States, and to develop the computational infrastructure needed to reach these goals. Membership in USQCD is open to all physicists based in the United States, and nearly all the high energy and nuclear physicists in the country working on the numerical study of QCD are members. Overall leadership of USQCD is vested in the Lattice QCD Executive Committee, whose members are the authors of this proposal. They have been working together since 1999 to organize the community, develop plans for the infrastructure, obtain funding to carry out these plans and oversee the implementation of them. The Executive Committee has appointed the Scientific Program Committee (SPC), which plays a major role in setting scientific priorities and allocating USQCD resources, as is described in Subsection 4.1. Members serve three year terms. The current ones are Tom Blum (U. Connecticut), Chris Dawson (U. Virginia), Frithjof Karsch (BNL), Andreas Kronfeld (FNAL, Chair), Colin Morningstar (Carnegie Mellon U.), John Negele (MIT), and Junko Shigemitsu (Ohio State). Additional information regarding the USQCD Collaboration can be found at its web site http://www.usqcd.org.

Funding of computational infrastructure by the DOE over the last seven years has played a critical role in maintaining a strong research program in lattice QCD in the United States. The SciDAC I grant, which ran from 2001 to 2006, and the SciDAC II grant, which began in 2006 and runs through 2011 have supported the development of community software, which is briefly described in Subsection 4.3. As indicated in the previous section, the SciDAC I grant funded the acquisition and operation of prototype clusters in the period 2001–2006, a separate grant funded the construction of the QCDOC in 2004-2005, and LQCD has funded the acquisition of clusters and the operation of all USQCD hardware since 2006.

4.1 Scientific Priorities and Hardware Allocation

The SPC plays the leading role in setting scientific priorities and allocating USQCD computational resources, activities which are closely intertwined. Once a year it issues a call for proposals for use of USQCD's dedicated computers. It also invites the authors of the proposals to identify projects that would be appropriate for the DOE's leadership class computers. Proposals are grouped into three classes. Type A proposals are for investigations of very large scale, which will require a substantial fraction of the available resources. They are expected to benefit the whole Collaboration by producing data, such as gauge configurations or quark propagators, that can be used by all, or by producing physics results listed among USQCD's strategic goals. Type B proposals are for investigations of medium to large scale, which will require a smaller amount of resources than Type A ones. Type B proposal are not required to share data or to work towards stated Collaboration goals, although if they do, it is a plus. Such proposals may be scientifically valuable even if not closely aligned with USQCD goals. Type C proposals are for exploratory calculations, such as those needed to develop and/or benchmark code, acquire expertise in the use of USQCD hardware, or to perform investigations of limited scope. Whereas Type A and B proposals must be made in writing in response to the SPC's yearly call for proposals, Type C proposals can be made at any time simply by contacting a designated person at the Laboratory where the hardware ones wishes to use is located. Type A projects that generate multi-purpose data are candidates to be included in proposals for time on leadership class machines if their proponents can demonstrate the need for such resources and the ability to use them effectively.

After the Type A and B proposals are submitted, the SPC makes a preliminary review of them, and organizes the yearly All Hands meeting of the Collaboration. Authors of the proposals are invited to present their plans, which are discussed by the Collaboration as a whole. Following that discussion the Scientific Program Committee allocates time on USQCD resources, and transmits to the Lattice QCD Executive Committee the priorities of the Collaboration for use of the leadership class machines. It is the responsibility of the Executive Committee to submit proposals for the use of these computers on behalf of the Collaboration. This process sets the priorities of USQCD on a yearly basis.

The process just outlined has a number of advantages. Research priorities are examined and, if necessary, reordered on a yearly basis. The Executive Committee appoints members of the SPC with the aim of assembling a highly respected group of lattice gauge theorists that is balanced among the major research areas pursued by USQCD. The SPC, with input from the community as a whole, is thus in an excellent position to carry out its responsibilities. We believe that by setting priorities and allocating resources with such broad community participation, we will best advance the science and make the most efficient use of the resources available to us. Although we believe that it is important to apply for the large resources offered by the Incite Program as a community in order to build a coherent research effort, we also want to leave room for innovative projects that might not be central to the aims of USQCD at the moment. We have therefore made it a point not to apply for the smaller resources offered by NERSC or the NSF as a community, leaving such proposals to individuals or groups within USQCD.

An important success of this broadly representative allocation process is the impact it has on younger physicists. In order to obtain tenure track faculty positions and eventually be promoted to tenure, they must gain visibility and demonstrate independence. This can be difficult because most lattice QCD calculations require large computational resources, which can be difficult for beginning scientists to obtain at national centers. Members of the SPC are able to realistically evaluate the prospects of beginning lattice gauge theorists, and they have made a concerted effort to solicit and fund proposals from younger members of the field. For example, during the last USQCD allocation process, three outstanding postdoctoral research associates were awarded 6% of the total resources on dedicated computers for their calculation of the $K^0-\overline{K}^0$ mixing parameter B_K.

4.2 International Collaborations and Cooperation

Lattice QCD is an international field with very strong programs in Germany, Italy, Japan, and the United Kingdom, and excellent work being done in many other countries. A variety of quark actions are currently being used. In addition to the anisotropic Clover, DWF and improved staggered actions widely used in the US, overlap, twisted mass and isotropic Clover actions are studied extensively in other countries. The use of different quark actions serves a vital scientific purpose, allowing for studies of the effectiveness different actions for different problems, and for cross-checks of important results.

Groups within USQCD have formed collaborations with colleagues in other countries who share physics objectives and make use of the same actions. Four of these collaborations are particularly noteworthy. The present USQCD effort using DWF quarks is explicitly an international effort begun with equal participation by the United States based RBC collaboration and Edinburgh/Southampton/Swansea members of the UKQCD collaboration. During the past year this activity has expanded to include the LHPC group in the US. The Fermilab Lattice, HPQCD and MILC Collaborations have worked together in various combinations to study heavy quark physics using improved staggered quarks. HPQCD includes physicists in both USQCD and UKQCD. The members of the RBC Collaboration studying QCD thermodynamics using the P4 staggered quark action have a a long term collaboration with physicists at the University of Bielefeld, Germany. The effort to understand the hadron spectrum using Clover quarks on anisotropic lattices has close ties with the Trinity College Dublin group. In addition to their scientific work, our international collaborators have made major contributions to the development of algorithms and software, and have provided very significant computational resources to the joint efforts.

USQCD plays an active role in the International Data Grid (ILDG) with four members on the ILDG Metadata and Middleware Working groups, who co-ordinate standards with the USQCD Software Committee. The ILDG provides the means for sharing gauge configurations and quark propagators on an even wider international level. The ILDG has developed standards for file format and content, and the middleware needed to archive and retrieve files. Further information regarding the ILDG can be found at http://www.usqcd.org/ildg/ Sharing of these valuable data sets enables the international lattice QCD community to maximize the science it can produce from the computational resources available. USQCD unquestionably gains significantly from its participation in this effort.

The resources we request are based on the requirements of the research program set out in Section 2. However, it may help to put them in perspective by comparing our current resources with those available for the study of lattice QCD in other countries. We do this in Table 11, where we show estimates of the computing resources available for the study of lattice QCD in the countries that are major participants in the field, as of October, 2007. The estimates for other countries were obtained by making inquiries of senior physicists in each of them, and translating their responses into our standard measure, the average of the sustained performance of the routines for computing DWF and Asqtad quark propagators. Approximately one-third of the United States resources labeled National Centers come from an allocation to the USQCD Collaboration by the DOE's INCITE Program, while the remainder comes from allocations to individuals or groups by NSF Centers and NERSC. Three computers located in the United States, but also not allocated by the USQCD Collaboration, are not shown in the table. One is a QCDOC located at the Riken Brookhaven Research Center, which was funded by the Riken Institute of Japan, and is used by the RBC Collaboration. The second is the NNSA BlueGene/L located at Lawrence Livermore National Laboratory, which is being used by the HotQCD and NPQCD collaborations, and the third is the BlueGene/L, New York Blue, which is being accessed by physicists at Columbia and BNL in its early user period. It is clear that without LQCD the United States would have been very hard pressed to maintain a world class research program in lattice QCD. Physicists in other countries also recognize the scientific opportunities in lattice QCD that will be available over the next several years, and are moving aggressively to obtain the computing resources necessary to capitalize on them. They are seeking resources similar to those we propose, and it is our understanding that major upgrades to those listed in Table 11 have been, or are about to be approved. Thus, for US physicists to effectively collaborate in this rapidly developing international environment, it is important that we have access to resources of the scale proposed here.

Country	Sustained Teraflop/s
Germany	10–15
Italy	5
Japan	14–18
United Kingdom	4–5
Unites States	
LQCD Project	12
National Centers	3
US Total	15

Table 11: Computing resources in sustained teraflop/s estimated to be available for the study of Lattice QCD in various countries, as of October, 2007.

4.3 SciDAC Software

Under our SciDAC I and SciDAC II grants we have developed software that enables us to write highly efficient and portable codes for the study of lattice QCD. Under SciDAC I a QCD Applications Programming Interface (QCD API) was created with the design goals of enabling users to quickly adapt codes to new architectures, easily develop new applications and incorporate new algorithms, and preserve their large investment in existing codes.

The QCD API has a layered structure which is implemented in a set of independent libraries. Level 1 provides the code that controls communications (QMP) and the basic linear algebra routines that run on individual compute cores (QLA). To obtain high efficiency much of this layer may have to be written in hardware specific assembly language and/or make use of low level communications calls. However versions exist in C and C++ using MPI for transparent portability of all application codes. Level 2 (QDP) contains data parallel operations that are built on QMP and QLA. QDP allows extensive overlapping of communication and computation in a single line of code. By making use of the QMP and QLA layers, the details of communications buffers, synchronization barriers, vectorization over multiple sites on each node, etc. are hidden from the user. Both C and C++ versions of QDP are in use. A very large fraction of the resources in any lattice QCD simulation go into a few computationally intensive subroutines, most notably the repeated inversion of the Dirac operator, a large sparse matrix. To obtain the level of efficiency at which we aim, it is necessary to optimize these subroutines for each architecture. Level 3 codes have been written in assembly language for the PowerPC processors used in the QCDOC and BlueGene computers, and with SSE instructions for AMD and Intel processors used in our clusters and in the Cray XT4. An I/O library (QIO) that adheres to the ILDG file standards enables parallel I/O and sharing of large data sets among members of USQCD and with the international community via the ILDG.

Under SciDAC II the QCD API is being used to enhance the performance of the two pre-existing publicly available codes written by members of USQCD: the Columbia Physics System (CPS) and the MILC codes. It is also being used to extend a third publicly available code, CHROMA, which was written entirely under the C++ version of the QCD API during SciDAC I. In addition, the QCD API is being optimized for BlueGene and Cray supercomputers, and for clusters based on multi-core processors. A QCD physics toolbox is being constructed which will contain sharable software building blocks for inclusion in application codes, performance analysis and visualization tools, and software for automation of physics work flow.

The software created under the SciDAC grants has greatly enhanced the effectiveness with which we are using the hardware resources available to us, and will continue to do so as more of the SciDAC II initiatives are completed. All of the software developed under the SciDAC grants is publicly available, and can be found at http://usqcd.jlab.org/usqcd-software.

5 Role of the Laboratories

The three participating laboratories, BNL, FNAL and TJNAF, have played a critical role in LQCD, and will do so again in LQCD-ext. They make available the outstanding staff who evaluate, procure and operate the dedicated hardware. They also provide exceptional user support. It would be very difficult, if not impossible, for us to obtain the services of such talented staff on our own. This approach is also highly cost effective, as our operations budget is only charged for the staff hours actually devoted to our effort. The laboratories also provide space, air conditioning and electricity for the hardware from their base budgets. FNAL and TJNAF give us access to their archival storage systems, with our only charge being for the tapes. We also benefit from the laboratories' excellent networking infrastructure and their computer security programs. We believe that these crucial contributions are only possible because three laboratories are involved. If the hardware were located at a single site, then we would very likely become a significant drain on that laboratory's resources, rather than the small perturbation we presently are.

All three laboratories have indicated that they will request funds from the DOE to purchase dedicated hardware for lattice QCD only through the national effort set out in this proposal. In LQCD, FNAL and TJNAF have acquired and operated highly cost effective commodity clusters for our community, and propose to continue to do so in LQCD-ext. BNL has operated the QCDOC for us, and will continue to do so through 2011. The hardware interests and expertise at BNL are now focused on the BlueGene line of supercomputers, and they have proposed to procure and support BlueGene hardware for us. Again, in any given year we will acquire the hardware that best advances our science.

The contributions of the laboratories go well beyond provision of staff and physical infrastructure. Each of them has an outstanding lattice QCD group. Among them, the three laboratory groups cover the major areas of research set out in Section 2. They serve as intellectual centers for work in these areas. A particularly important role of the laboratory groups is to stimulate interactions between members of the lattice QCD community and experimentalists. Indeed, the committee that reviewed LQCD in the spring of 2007 stated in its report that "*Contact with experiment is strong, partly because lattice QCD has computational and human resources spread over three national laboratories.*" Some recent highlights are as follows. The *BaBar/Lattice QCD Workshop*, which took place at SLAC on September 16, 2006, brought together members of USQCD working on the calculation of weak interaction matrix element with members of the BaBar experiment. It was so successful that a second workshop, *Lattice QCD Meets Experiment 2007*, was organized at FNAL on December 10-11, 2007. It was expanded to include experimentalists from Babar, CDF, CLEO-c, and D0. A series of workshops has been held at BNL to bring together experimentalists working at RHIC and those studying QCD thermodynamics on the lattice. The latest of these was titled *Can We Discover the QCD Critical Point at RHIC*, which took place on March 9-10, 2006. The next one, *Understanding QGP through Spectral Functions and Euclidean Correlators*, is scheduled to take place at BNL on April 23-25, 2008. It will be followed by one titled *Critical Point and Onset of Deconfinement*, scheduled for March 16–20, 2009 in conjunction with a planned low-energy scan at RHIC in FY 2010 to search for this critical point. In the area of hadronic and nuclear physics, the workshop on *Synergy Between Experiment and Lattice QCD in Exploring Hadron Structure* was held at the Institute for Nuclear Theory on April 24-25, 2006 with joint sponsorship by INT and TJNAF, and organizational work by members of USQCD. Furthermore, TJNAF members of USQCD are particularly active in interacting with experimentalists, and are even involved in planning experiments.

6 Management Plan

The 2007 LQCD Computing Project Review Report stated that "*The organizational structure appears to be optimal in providing a platform for large-scale computing to the US lattice QCD community.*" We agree, and therefore propose that the management structure now in place for LQCD be used for LQCD-ext. William Boroski (FNAL), the Project Manager, has overall responsibility for the effort. He is responsible for assuring that the project is well defined via a work breakdown structure and well tracked via milestones. He is the key interface to the DOE for financial matters, reporting and reviews. He is assisted by the Associated Project Manager, Bakul Banerjee (FNAL). She is responsible for maintaining the work breakdown structure and other controlled documents. She also tracks expenditures and progress in achieving milestones. Each participating laboratory has a Site Manager: Eric Blum at BNL, Donald Holmgren and Amitoj Singh at FNAL, and Chip Watson at TJNAF. The Site Managers are responsible for hardware selection, procurement, deployment and operation at their sites consistent with the overall plan. They are also responsible for site operations and user support. The Chair of the LQCD Executive Committee, Robert Sugar, serves as the Scientific Spokesperson for the effort. He is the principal point of contact with the DOE on scientific matters, and liaison between the Executive Committee and the Project Manager, relating the Executive Committee's priorities to the Project Manager, and the Program Manager's progress reports to the Executive Committee. The Project Manager, Associate Project Manager, Site Managers and Scientific Spokesperson meet approximately every other week by conference call to discuss major issues.

The Change Control Board (CCB) evaluates the feasibility, cost and impact of proposed changes to the project which result in more than a minimal cost or schedule change. Its role is to assure that proposed changes are managed with the primary focus on the advancement of the scientific goals of the project. The members are the Project Manager, the Chair of the LQCD Executive Committee, who serves as chair of the CCB, the FNAL Computing Division Head, Victoria White, the TJNAF Chief Information Officer, Roy Whitney, the BNL Information Technology Division Director, Tom Schlagel, and a member of the USQCD Collaboration selected by the Executive Committee, Steven Gottlieb.

Appendix

Revision history. This document differs from the version submitted in January, 2008 in the following ways. The project names LQCD and LQCD-ext for the original project and its extension have replaced the names LQCD-I and LQCD-II that were used in the 2008 proposal. The budget numbers used in Table 10, in the paragraph preceeding Table 10, and in the Introduction have been revised to reflect current anticipations. The anticipated dedicated hardware in Table 8 has been revised to reflect these budget numbers, and the paragraph preceeding Table 8 has been revised accordingly.

References

[1] More extensive discussions of the four areas of research addressed in this section can be found in white papers located on the USQCD web page http://www.usqcd.org/collaboration.html#2007Whitepapers.

[2] C. T. H. Davies *et al.* [HPQCD, UKQCD, MILC, and Fermilab Collaborations], Phys. Rev. Lett. **92**, 022001 (2004) [arXiv:hep-lat/0304004].

[3] C. Aubin *et al.* [MILC Collaboration], Phys. Rev. D **70**, 114501 (2004) [arXiv:hep-lat/0407028].

[4] Q. Mason *et al.* [HPQCD Collaboration], Phys. Rev. Lett. **95**, 052002 (2005) [arXiv:hep-lat/0503005].

[5] These predictions are summarized in A. S. Kronfeld [Fermilab Lattice Collaboration Collaboration], arXiv:hep-lat/0607011.

[6] G.P. Lepage, private communication.

[7] C. Bernard *et al.* [MILC collaboration], arXiv:hep-lat/0611024.

[8] M. Nobes and H. Trottier, PoS **LAT2005**, 209 (2006) [arXiv:hep-lat/0509128].

[9] A. Gray, I. Allison, C. T. H. Davies, E. Dalgic, G. P. Lepage, J. Shigemitsu and M. Wingate, Phys. Rev. D **72**, 094507 (2005) [arXiv:hep-lat/0507013].

[10] USQCD workshop “Lattice QCD meets Experiment 2007”, Fermilab, Dec. 2007, http://usqcd.fnal.gov/lattice-experiment2007/schedule.html.

[11] USQCD white paper “Fundamental Parameters from Future Lattice Calculations”, Jan. 2007, http://www.usqcd.org/documents/fundamental.pdf.

[12] “UTfit” web site, http://utfit.roma1.infn.it/. Similar plots are available on the the CKM fitter web site: http://www.slac.stanford.edu/xorg/ckmfitter/ckm_welcome.html.

[13] D.J. Antonio *et al.*, [RBC and UKQCD Collaborations], arXiv:hep-ph/0702042.

[14] E. Dalgic *et al.*, [HPQCD collaboration] Phys. Rev. D **76**, 011501 (2007) [arXiv:hep-lat/0610104].

[15] M. Okamoto, PoS **LAT2005**, 013 (2006) [arXiv:hep-lat/0510113].

[16] V. Fanti *et al.* [NA48 Collaboration], Phys. Lett. B **465**, 335 (1999) [arXiv:hep-ex/9909022].

[17] A. Alavi-Harati *et al.* [KTeV Collaboration], Phys. Rev. Lett. **83**, 22 (1999) [arXiv:hep-ex/9905060].

[18] L. Lellouch and M. Luscher, Commun. Math. Phys. **219**, 31 (2001) [arXiv:hep-lat/0003023].

[19] S. Li, R. D. Mawhinney [RBC and UKQCD Collaborations], arXiv:0710.3414 [hep-lat].

[20] P. Boucaud, V. Gimenez, C. J. D. Lin, V. Lubicz, G. Martinelli, M. Papinutto and C. T. Sachrajda, Nucl. Phys. B **721**, 175 (2005) [arXiv:hep-lat/0412029].

[21] C. J. D. Lin, G. Martinelli, E. Pallante, C. T. Sachrajda and G. Villadoro, Phys. Lett. B **553**, 229 (2003) [arXiv:hep-lat/0211043].

[22] C. H. Kim and C. Sachrajda, arXiv:0710.2519 [hep-lat].

[23] C. H. Kim, Nucl. Phys. Proc. Suppl. **140**, 381 (2005).

[24] C. H. Kim and N. H. Christ, Nucl. Phys. Proc. Suppl. **119**, 365 (2003) [arXiv:hep-lat/0210003].

[25] T. Yamazaki [RIKEN-BNL-Columbia Collaboration], PoS **LAT2006**, 100 (2006) [arXiv:hep-lat/0610051].

[26] One Tflop-Year is defined to be the number of floating points operations produced in one year by a computer that sustains one teraflop/s.

[27] Throughout this proposal, performance is measured as the average of that sustained by the sparse matrix inversion routines for computing the quark propagators for the Domain Wall and Improved Staggered (Asqtad) quark actions under production conditions.

[28] E. Follana *et al.* [HPQCD Collaboration], Phys. Rev. D **75**, 054502 (2007) [arXiv:hep-lat/0610092].

[29] P.M. Vranas, arXiv:hep-lat/0001006.

[30] H. Matsufuru [JLQCD Collaboration], arXiv:0710.4225 [hep-lat].

[31] R.C. Brower, H. Neff and K. Orginos, Nucl. Phys. Proc. Suppl. **153**, 191 (2006) [arXiv:hep-lat/0511031].

[32] C. Bernard, *et al.* [MILC Collaboration], Phys. Rev. **D75**, 094505 (2007) [arXiv:hep-lat/0611031].

[33] M. Cheng, *et al.*, arXiv:0710.0354 [hep-lat].

[34] S. Ejiri, F. Karsch, E. Laermann and C. Schmidt, Phys. Rev. **D73**, 054506 (2006) [arXiv:hep-lat/0512040].

[35] E. Follana, *et al.* [HPQCD Collaboration], Phys. Rev. **D75**, 094502 (2007) [arXiv:hep-lat/0610092].

[36] S. Ejiri, *et al.*, Prog. Theor. Phys. Suppl. **153**, 118 (2004) [arXiv:hep-lat/0312006].

[37] G. Aarts, C. Allton, M. Oktay, M. Peardon and J.-I. Sullerud, Phys. Rev. **D76**, 094513 (2007) arXiv:0705.2198 [hep-lat].

[38] M. Asakawa and T. Hatsuda, Phys. Rev. Lett. **92**, 012001 (2004) [arXiv:hep-lat/0308034].

[39] A. Jakovac, P. Petreczky, K. Petrov and A. Velysky, Phys. Rev. **D75**, 014506 (2007) [arXiv:hep-lat/0611017].

[40] F. Karsch, D. Kharzeev and K. Tuchin, arXiv:0711.0914 [hep-ph].

[41] H. Meyer, arXiv:0710.3717 [hep-lat].

[42] P. de Forcrand and O. Philipsen, JHEP **01**, 077 (2007) [arXiv:hep-lat/0607017].

[43] Performance measure in Nuclear Physics adopted by the Nuclear Science Advisory Committee are posted at the url:
http://www.sc.doe.gov/henp/np/nsac/docs/nsac_report_performance_measures.pdf

[44] D. Dolgov *et al.* [LHPC collaboration], Phys. Rev. D **66**, 034506 (2002) [arXiv:hep-lat/0201021].

[45] The LHC Collaboration: P. Hagler, *et al.*, Phys. Rev **D68**, 034505 (2003).

[46] The LHC Collaboration: P. Hagler, *et al.*, Phys. Rev. Lett. **93**, 112001 (2004).

[47] C. Alexandrou, P. de Forcrand, Th. Lippert, H. Neff, J. W. Negele, K. Schilling, W. Schroers and A. Tsapalis, Phys. Rev **D69**, 114506 (2004).

[48] C. Alexandrou, P. de Forcrand, H. Neff, J. W. Negele, W. Schroers and A. Tsapalis, Phys. Rev. Lett. **94**, 021601 (2005).

[49] R. G. Edwards *et al.* [LHPC Collaboration], Phys. Rev. Lett. **96**, 052001 (2006) [arXiv:hep-lat/0510062].

[50] D. B. Renner *et al.* [LHPC Collaboration], J. Phys. Conf. Ser. **46**, 152 (2006) [arXiv:hep-lat/0607008].

[51] D. B. Renner *et al.* [LHPC Collaboration], arXiv:0710.1373 [hep-lat].

[52] Ph. Hagler *et al.* [LHPC Collaborations], arXiv:0705.4295 [hep-lat].

[53] R. G. Edwards *et al.*, arXiv:hep-lat/0610007.

[54] X. D. Ji, Phys. Rev. Lett. **78**, 610 (1997) [arXiv:hep-ph/9603249].

[55] C. Alexandrou, T. Leontiou, J. W. Negele and A. Tsapalis, arXiv:hep-lat/0607030.

[56] C. Alexandrou, G. Koutsou, T. Leontiou, J. W. Negele and A. Tsapalis, Phys. Rev. D **76**, 094511 (2007).

[57] C. Alexandrou, G. Koutsou, H. Neff, J. W. Negele, W. Schroers and A. Tsapalis, arXiv:0710.4621 [hep-lat].

[58] S. R. Beane, K. Orginos and M. J. Savage, arXiv:hep-lat/0605014.

[59] S. Ohta and K. Orginos [RBCK Collaboration], Nucl. Phys. Proc. Suppl. **140**, 396 (2005) [arXiv:hep-lat/0411008].

[60] H. W. Lin and S. Ohta, PoS **LAT2006**, 118 (2006) [arXiv:hep-lat/0610028].

[61] T. Yamazaki and S. Ohta [RBC Collaboration], arXiv:0710.0422 [hep-lat].

[62] C. J. Morningstar and M. J. Peardon, Phys. Rev. D **60**, 034509 (1999) [arXiv:hep-lat/9901004].

[63] Y. Chen *et al.*, Phys. Rev. D **73**, 014516 (2006) [arXiv:hep-lat/0510074].

[64] The LHPC Collaboration: S. Basak, *et al.*, Phys. Rev **D72**, 094506 (2005).

[65] The LHPC Collaboration: S. Basak, *et al.*, Phys. Rev **D72**, 074501 (2005).

[66] A. Lichtl, arXiv:hep-lat/0609019.

[67] S. Basak *et al.*, Phys. Rev. D **76**, 074504 (2007) [arXiv:0709.0008 [hep-lat]].

[68] J. J. Dudek, R. G. Edwards, N. Mathur and D. G. Richards, arXiv:0707.4162 [hep-lat].

[69] J. J. Dudek, R. G. Edwards and D. G. Richards, Phys. Rev. D **73**, 074507 (2006) [arXiv:hep-ph/0601137].

[70] J. J. Dudek and R. G. Edwards, Phys. Rev. Lett. **97**, 172001 (2006) [arXiv:hep-ph/0607140].

[71] S. R. Beane, P. F. Bedaque, A. Parreno and M. J. Savage, Phys. Lett. B **585**, 106 (2004) [arXiv:hep-lat/0312004].

[72] S. R. Beane, P. F. Bedaque, K. Orginos and M. J. Savage [NPLQCD Collaboration], Phys. Rev. D **73**, 054503 (2006) [arXiv:hep-lat/0506013].

[73] S. R. Beane, T. C. Luu, K. Orginos, A. Parreno, M. J. Savage, A. Torok and A. Walker-Loud, arXiv:0706.3026 [hep-lat].

[74] S. R. Beane, P. F. Bedaque, K. Orginos and M. J. Savage, Phys. Rev. Lett. **97**, 012001 (2006) [arXiv:hep-lat/0602010].

[75] S. R. Beane, P. F. Bedaque, T. C. Luu, K. Orginos, E. Pallante, A. Parreno and M. J. Savage, arXiv:hep-lat/0612026.

[76] S. R. Beane, P. F. Bedaque, T. C. Luu, K. Orginos, E. Pallante, A. Parreno and M. J. Savage, arXiv:hep-lat/0607036.

[77] H. B. Meyer and J. W. Negele, arXiv:0707.3225 [hep-lat].

[78] J. Foley, K. Jimmy Juge, A. O'Cais, M. Peardon, S. M. Ryan and J. I. Skullerud, Comput. Phys. Commun. **172**, 145 (2005) [arXiv:hep-lat/0505023].

[79] M. Luscher, Commun. Math. Phys. **105**, 153 (1986).

[80] S. Aoki *et al.* [CP-PACS Collaboration], Phys. Rev. D **76**, 094506 (2007) [arXiv:0708.3705 [hep-lat]].

[81] J. Kuti, L. Lin and Y. Shen, Phys. Rev. Lett. **61**, 678 (1988).

[82] M. Luscher and P. Weisz, Phys. Lett. B **212**, 472 (1988).

[83] W. Bock, A. K. De, K. Jansen, J. Jersak, T. Neuhaus and J. Smit, Nucl. Phys. B **344**, 207 (1990); I. H. Lee, J. Shigemitsu and R. E. Shrock, ibid Nucl. Phys. B **334**, 265 (1990); L. Lin, I. Montvay, H. Wittig and G. Munster, Nucl. Phys. B **355**, 511 (1991); K. Jansen, J. Kuti and C. Liu, Phys. Lett. B **309**, 119 (1993); ibid Phys. Lett. B **309**, 127 (1993), K. Holland and J. Kuti, Nucl. Phys. Proc. Suppl. **129**, 765 (2004); K. Holland, Nucl. Phys. Proc. Suppl. **140**, 155 (2005).

[84] Z. Fodor, K. Holland, J. Kuti, D. Nogradi, C. Schroeder, PoS(LATTICE 2007)056; P. Gerhold and K. Jansen, JHEP **0709**, 041 (2007).

[85] Fleming, George Tamminga and Kogut, John B. and Vranas, Pavlos M., Phys. Rev. D64 034510 (2001).

[86] D. B. Kaplan, Nucl. Phys. Proc. Suppl., 129,109, (2004).

[87] J. Giedt, Int. J. Mod. Phys. A21, 3039 (2006).

[88] S. Catterall, JHEP, 11, 006 (2004).

[89] S. Catterall, JHEP 03,032 (2006).

[90] S. Catterall and T. Wiseman, arXiv:0706.3518 [hep-lat] (2007).

[91] T. Banks and A. Zaks, Nucl. Phys. B196, 189 (1982).

[92] T. Appelquist, G. T. Fleming and E. T. Neil, arXiv:0712.0609 [hep-ph] (2007).

[93] S. Catteralland F. Sannino, Phys. Rev. D76, 034504 (2007).

[94] H. Georgi, Phys. Rev. Lett. 98,221601 (2007) .

[95] R. Babich, R. C. Brower, M. Clark, G. Fleming, J. Osborn and C. Rebbi, PoS (LATTICE 2007) 139.

[96] E. Baltz, M. Battaglia, M. Peskin and T. Wizansky, Phys. Rev. D74, 03521 (2006).

[97] M. Golterman, Nucl. Phys. Proc. Suppl. **94**, 189 (2001); M. Luscher, arXiv:hep-th/0102028 (2001); M. Creutz, Rev. Mod. Phys. **73**, 119 (2001).

[98] M. A. Clark and A. D. Kennedy, Phys. Rev. Lett. **98**, 051601 (2007) [arXiv:hep-lat/0608015] and Phys. Rev. D **75**, 011502 (2007) [arXiv:hep-lat/0610047].

[99] Pavlos Vranos, Journal of Physics, Conference Series, **78**, 012080 (2007).